

DYNAMICS-ADAPTED CONE KERNELS

DIMITRIOS GIANNAKIS*

Abstract. We present a family of kernels for analysis of data generated by dynamical systems. These so-called cone kernels feature an explicit dependence on the dynamical vector field operating in the phase-space manifold, estimated empirically through finite-differences of time-ordered data samples. In particular, cone kernels assign strong affinity to pairs of samples whose relative displacement vector lies within a narrow cone aligned with the dynamical vector field. As a result, in a suitable asymptotic limit, the associated diffusion operator generates diffusions along the dynamical flow, and is invariant under a weakly restrictive class of transformations of the data, which includes conformal transformations. Moreover, the corresponding Dirichlet form is governed by the directional derivative of functions along the dynamical vector field. The latter feature is metric-independent. The diffusion eigenfunctions obtained via cone kernels are therefore adapted to the dynamics in that they vary predominantly in directions transverse to the flow. We demonstrate the utility of cone kernels in nonlinear flows on the 2-torus and North Pacific sea surface temperature data generated by a comprehensive climate model.

Key words. kernel methods, diffusion operators, eigenfunctions, manifold embedding, vector field, delay coordinates

AMS subject classifications. 37M10, 37N10

1. Introduction. Large-scale datasets generated by dynamical systems arise in a diverse range of disciplines in science and engineering, including fluid dynamics [32, 50], materials science [36, 35], molecular dynamics [20, 43], and geophysics [40, 22]. A major challenge in these domains is to utilize the vast amount of data that is being collected by observational networks or output by large-scale numerical models to understand the operating physical processes, and make inferences about aspects of the system which are not accessible to observation. For instance, in climate atmosphere ocean science (CAOS) the dynamics takes place in an infinite-dimensional phase space where the coupled nonlinear partial differential equations for fluid flow and thermodynamics are defined, and the observed data correspond to functions of that phase space, such as temperature or circulation measured over a geographical region of interest. There exists a strong need for data analysis algorithms to extract and create reduced representations of the large-scale coherent patterns which are an outcome of these dynamics, including the El Niño Southern Oscillation (ENSO) in the ocean [53] and the Madden-Julian Oscillation in the atmosphere [39]. Advances in the scientific understanding and forecasting capability of these phenomena have potentially high socioeconomic impact.

Despite the high phase space dimension of many systems of interest, their dynamics evolve asymptotically (at long times) on low-dimensional submanifolds of phase space (attractors) [22, 34, 2]. It is therefore natural to exploit

*Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences, New York University, New York, New York. (dimitris@cims.nyu.edu).

topological and geometrical aspects of these submanifolds to design data analysis algorithms for data reduction and decomposition, function learning, and other important problems. Here a major challenge stems from the fact that the geometry of the attractive manifolds is nonlinear, meaning that linear variance-optimizing algorithms such as principal components analysis (PCA) [3, 32] or singular spectrum analysis (SSA) [11, 54, 24] are likely to perform suboptimally. Indeed, it has been documented in the literature (e.g., [4, 18]) that the dynamically significant modes in nonlinear systems are not necessarily those carrying high variance. Moreover, in practical applications involving finite sample counts and short observation intervals, the data manifolds are inherently discrete, and represent only coarse-grained aspects of the underlying geometry at the infinite-sample limit.

Kernel methods have been extensively used as alternatives to classical linear algorithms to take advantage of nonlinear geometric structures of data. Prominent applications include dimension reduction and feature extraction [48, 5, 14, 15, 6, 33], learning and regularization of scalar or vector-valued functions [7, 42, 13], and out-of-sample extension of these functions [16]. Here, a common theme is that suitably constructed kernels, i.e., functions measuring a notion of pairwise similarity between data points, lead naturally to diffusion operators which are closely related to the manifold structure of the data. More specifically, because every elliptical diffusion operator induces a Riemannian metric tensor on the data manifold [46, 23], using that operator (or, as is frequently the case, its eigenfunctions) for tasks such as data representation and function learning is tantamount to performing these tasks in a manner compatible with the induced Riemannian geometry. Applied to datasets with nonlinear manifold structure, such as those generated by complex dynamical systems, this approach has been found to yield more efficient algorithms and meaningful results than linear PCA-type algorithms (e.g., [37] and references therein).

A major advantage of constructing diffusion operators through kernels is that kernels are defined in the ambient data space, and thus can be used to control the induced Riemannian metric [9] exploiting features of data space which are available for the problem at hand. In the context of dynamical system data, an important feature which is not present in general point clouds is that the samples occur with a time ordering which is the outcome of dynamical flow in phase space. Equivalently, we think of these datasets as being associated with a nowhere-vanishing vector field v on the attractor, which is intrinsic to the system under study in the sense that it does not depend on how the data manifold is embedded in ambient space. These observations provide a motivation to seek aspects of v which are empirically accessible in ambient data space, and incorporate them in dynamics-adapted kernels.

Efforts in that direction have been made in a series of papers on so-called nonlinear Laplacian spectral analysis (NLSA) algorithms for decomposition of spatiotemporal data [25, 28, 29], and in independent work by Berry and

collaborators [10, 9]. In both of these works, dynamical flow is incorporated in kernels by first embedding the data in a higher-dimensional space (hereafter called embedding space) through Takens' method of delays [44, 51, 11, 47]. Because each point in embedding space corresponds to a segment of dynamical evolution observed over a temporal window, distances in that space depend on the dynamical flow generating the data. This means in turn that the induced Riemannian metric associated with a kernel formulated in embedding space will depend (albeit indirectly) on the dynamical system generating the data.

In NLSA algorithms, the embedding-space distances are scaled by a factor which is proportional to the distance between temporal nearest neighbors. As will be made precise below, the displacement vector between temporal nearest neighbors is a representation of the dynamical vector field, pushed forward to embedding space. Therefore, in addition to time-lagged embedding, the NLSA kernel also depends directly on the dynamical vector field through its norm. Qualitatively, the result of this dynamics-dependent scaling factor is to assign higher weights to transitory states with large phase space velocity. This feature was found to be particularly useful in extracting dynamically significant modes in systems with metastability [28]. Moreover, because of the “non-dimensionalization” produced by the scaling factor, the NLSA kernel can be naturally extended to process multivariate datasets consisting of components with different physical units [12]. Yet, in spite of these attractive features, theoretical understanding of the role of the scaling factor in the induced Riemannian metric has so far been lacking.

Building on the existing work in [25, 28, 29], in this paper we introduce a one-parameter family of kernels which, besides the norm of the dynamical vector field also depend on the angle between that vector field and the displacement vector between the points at which the kernel is evaluated. In particular, the new kernels assign higher weight to pairs of points whose relative displacement lies within a cone with axis parallel to the dynamical vector field. For this reason we refer to the new kernels as cone kernels. The angular influence is controlled by a parameter $\zeta \in [0, 1)$, such that the existing NLSA kernel with no angular dependence occurs as the special case $\zeta = 0$.

Here, however, our main focus is the limiting behavior $\zeta \rightarrow 1$, where the angular influence is maximal. In that limit, the induced Riemannian metric becomes degenerate, and diffusion takes place along the integral curves of the dynamical vector field v (in the sense of [23]). Because the dynamical integral curves are intrinsic to the system (i.e., do not depend on the ambient-space induced metric), the along- v property of the $\zeta \rightarrow 1$ diffusion operator is invariant under arbitrary diffeomorphisms of the data manifold. Moreover, the associated Dirichlet form depends on the directional derivative of functions along v , which is also metric independent. In fact, the only dependence of the limit operator on the ambient-space metric is through a ratio involving its volume form and the norm of v . By virtue of this property, as $\zeta \rightarrow 1$, the diffusion operators constructed from cone kernels become invariant under

a weakly restrictive class of transformations of the data, which includes conformal transformations as a special case. In addition, due to the structure of the Dirichlet form, the associated diffusion eigenfunctions are expected to vary predominantly along directions transverse to v . These eigenfunctions can therefore be employed to carry out tasks such as dimension reduction and feature extraction in a dynamics-adapted manner.

We demonstrate the utility of cone kernels in two applications involving analytically solvable nonlinear flows on the 2-torus and dynamical evolution of sea surface temperature (SST) in the North Pacific ocean in a comprehensive climate model. In the torus application, we explicitly demonstrate the adaptation of the eigenfunctions of the associated diffusion operator to the dynamical flow, and invariance of these eigenfunctions under a non-conformal deformation of the torus. The climate model results suggest that the cone-kernel eigenfunctions produce more efficient (i.e., using fewer basis functions) representation of the temporal variability compared to the existing NLSA kernel, while achieving better timescale separation.

The plan of this paper is as follows. In section 2, we lay out the notation, and develop the cone kernel formulation. In section 3, we study the behavior of the induced metric and the associated diffusion operator. We present the applications to torus flows and climate model data in sections 4 and 5, respectively, and conclude in section 6. Technical results and proofs are included in an appendix. A movie illustrating the time evolution of SST filtered by diffusion eigenfunctions is provided as supplemental material. A Matlab code used to generate the numerical results in sections 4 and 5 is available upon request from the author.

2. Formulation of cone kernels. We consider a scenario where the dynamics is described by a deterministic flow $\Phi_t : \mathcal{F} \mapsto \mathcal{F}$ operating in a phase space \mathcal{F} , and evolving on a smooth (of class C^1), compact m -dimensional attractor $\mathcal{M} \subseteq \mathcal{F}$ without boundary. Moreover, observations are taken uniformly in time with a timestep $\delta t > 0$ on the attractor via a C^2 vector-valued function $F : \mathcal{F} \mapsto \mathbb{R}^n$, leading to a dataset


$$(2.1) \quad X = \{X_1, \dots, X_s\}, \quad \text{with} \quad X_i = F(a_i), \quad a_i = \Phi_{t_i} a_0, \quad t_i = i \delta t, \quad a_0 \in \mathcal{M}.$$

See Figure 1 for an illustration. We refer to \mathbb{R}^n interchangeably as data space or ambient space, and consider that it is equipped with an inner product (\cdot, \cdot) .

Throughout, we assume that Φ_t is sufficiently smooth so that the dynamical vector field v induced on \mathcal{M} , defined through

$$(2.2) \quad v(f) = \lim_{t \rightarrow 0} (f(\Phi_t a) - f(a))/t, \quad \text{with} \quad a \in \mathcal{M}, \quad f \in C^1 \mathcal{M},$$

is at least C^1 . Moreover, without loss of generality, we assume that F is one-to-one, and has full rank on the tangent spaces of \mathcal{M} ; i.e., F is an embedding of \mathcal{M} into \mathbb{R}^n (if F is not an embedding, it is generically possible to construct an



embedding through delay-coordinate maps [51, 47]). Under these conditions, the observation map generates a vector field V on $F(\mathcal{M})$ through its derivative DF , viz. $V = DF v$. We denote by g the Riemannian metric induced on \mathcal{M} by pulling back the ambient-space inner product, giving

$$(2.3) \quad q(u, u') = (DF u, DF u')$$

Remark 2.1. Even though we do not attempt to extend the analysis presented here to non-smooth manifolds, we note work of Sauer et al. [47], who prove embedding theorems for dynamical systems with fractal attractors. Moreover, in section 3.2 we pass to an intrinsically discrete formulation, where the existence of an underlying smooth continuous theory is not required.

Next, consider how linear combinations of X_i can be used to construct finite-difference (FD) approximations of V . In particular, let δ_p be a p -th order FD operator for the first derivative of a function,

$$(2.4) \quad \frac{df}{dt} = \frac{\delta_p f}{\delta t} + O(\delta t^p).$$

We then have the following Lemma.

LEMMA 2.1. *The vector in data space $\xi_i = \delta_p X_i$ corresponds to an $O(\delta t^p)$ -accurate FD approximation of the pushforward of the dynamical vector field v evaluated at $a_i \in \mathcal{M}$, in the sense that*

$$(2.5) \quad DF v|_{a_i} = \xi_i/\delta t + O(\delta t^p).$$

Moreover, the data space norm $\|\xi_i\| = \sqrt{(\xi_i, \xi_i)}$ provides an $O(\delta t^p)$ approximation to the norm of $v|_{a_i}$ with respect to the induced metric, namely

$$(2.6) \quad \|v\|_{q,a_i} := \sqrt{g(v,v)}|_{a_i} = \|\xi_i\|/\delta t + O(\delta t^p).$$

A proof of this Lemma is included in Appendix A.1. Hereafter, we nominally work with a central FD scheme,

$$(2.7) \quad \delta_p X_i = \sum_{j=-p}^p w_j X_{i+j},$$

where w_{-p}, \dots, w_p are standard FD weights for central schemes (e.g., [38]). However, the asymptotic analysis in section 3 only depends on the FD accuracy, and also applies, e.g., for backward and forward schemes.

Recall now that a kernel is a symmetric function $K : \mathcal{M} \times \mathcal{M} \mapsto \mathbb{R}_+$ which maps pairs of states in \mathcal{M} to a positive number, and, in practical applications, depends only on quantities observed in data space through F . A standard choice in this context is the isotropic Gaussian kernel [5, 14, 15, 6],

$$(2.8) \quad \bar{K}_\epsilon(a_i, a_j) = \exp(-\|\omega_{ij}\|^2/\epsilon), \quad \text{with} \quad \omega_{ij} = X_j - X_i = F(a_j) - F(a_i).$$

Here, ϵ and ω_{ij} are a positive parameter and the displacement vector between data samples X_i and X_j , respectively.

Having established the association between ξ_i and v_i in Lemma 2.1, we seek to modify \bar{K} to incorporate information about the dynamical vector field through (i) its norm $\|v\|_{g, a_i}$ with respect to the induced metric estimated via (2.6); (ii) its angle relative to ω_{ij} , estimated via

$$\cos \theta_i = \frac{(\xi_i, \omega_{ij})}{\|\xi_i\| \|\omega_{ij}\|}.$$

Specifically, introducing a parameter $\zeta \in [0, 1)$, we define

$$(2.9) \quad K_{\delta t, \zeta}(a_i, a_j) = \exp \left(-\frac{\|\omega_{ij}\|^2}{\|\xi_i\| \|\xi_j\|} [(1 - \zeta \cos^2 \theta_i)(1 - \zeta \cos^2 \theta_j)]^{1/2} \right).$$

For $\zeta = 0$, the kernel in (2.9) reduces to the locally scaled kernel employed in NLSA algorithms [28, 29], which features no dependence on the angle between ω_{ij} and the phase space velocity vectors, ξ_i and ξ_j . On the other hand, as ζ approaches 1, $K_{\delta t, \zeta}$ assigns higher affinity to data samples whose relative displacement vector is aligned with either ξ_i and/or ξ_j . For this reason, we term this two-parameter family of kernels cone kernels. For the remainder of this section we discuss certain properties of cone kernels which will be useful for the asymptotic analysis in section 3.

2.1. Qualitative features of cone kernels. As is evident from the structure of (2.9), the scaling of the pairwise distances $\|\omega_{ij}\|$ by $\|\xi_i\| \|\xi_j\|$ results in greater similarity being ascribed to transitory states characterized by large norm of the dynamical vector field. In [28], this feature was found to be crucial for successful dimensional reduction of a dynamical system with chaotic metastability. In section 3, we give a geometrical interpretation of the scaling factor showing that the metric tensor induced on \mathcal{M} through $K_{\delta t, \zeta}$ is invariant under conformal transformations of the data as a result of this scaling.

In addition, cone kernels with $\zeta \approx 1$ provide superior discrimination by assigning greater similarity to those sample pairs whose relative displacement vector is aligned with the dynamical flow. More specifically, given two distinct

data samples a_i and a_j with $\omega_{ij} \neq 0$, the ratio

$$\frac{K_{\delta t,1}(a_i, a_j)}{K_{\delta t,0}(a_i, a_j)} = \exp \left(\frac{\|\omega_{ij}\|^2 (1 - C_{ij})}{\|\xi_i\| \|\xi_j\|} \right), \quad C_{ij} = [(1 - \cos^2 \theta_i)(1 - \cos^2 \theta_j)]^{1/2} \leq 1$$

grows exponentially as $\delta t \rightarrow 0$ whenever $\cos \theta_i$ and $\cos \theta_j$ are both not equal to unity. An outcome of this property is that the corresponding diffusion operator Δ_ζ generates diffusions along the integral curves of the dynamical vector field v [23] in the sense that, asymptotically as $\zeta \rightarrow 1$, $\Delta_\zeta f$ vanishes whenever the gradient of f is parallel to v . Because these curves are intrinsic to the dynamical system generating the data, this property does not depend on the observation function F and the associated induced metric in (2.3), so long as F meets the conditions of a manifold embedding [47].

We also mention the utility of cone kernels in situations where F is a composite map, $F : \mathcal{M} \mapsto \mathbb{R}^{n_1} \oplus \mathbb{R}^{n_2}$, such that $F(a) = (F_1(a), F_2(a))$ where both F_1 and F_2 are embeddings. This scenario arises in practice when one has access to multivariate observations with distinct physical units, but there is no natural way of choosing a norm for the product space $\mathbb{R}^{n_1} \oplus \mathbb{R}^{n_2}$. Because the ratio $\|\omega_{ij}\|^2 / \|\xi_i\| \|\xi_j\|$ is invariant under scaling of the data by a constant (including change of units), cone kernels computed individually for F_1 and F_2 can be combined into a single product kernel without having to introduce additional scaling parameters. A climate science application of this technique can be found in [12].

2.2. Local behavior at the basepoint. The metric tensor induced at a reference point a on the data manifold by an exponentially decaying kernel depends strongly on the local rate of decay of the kernel at a . As established by Berry [9], in many cases of interest, including the isotropic Gaussian kernel in (2.8) and the cone kernel in (2.9), that rate is controlled to leading order by the Hessian (second derivative) matrix of the kernel with respect to a coordinate patch covering the reference point, as we now discuss.

Fixing a basepoint $a \in \mathcal{M}$, consider $K_{\delta t, \zeta}(a, a')$ for a' lying in an exponential neighborhood of a :

$$a' = \exp_a u, \quad u \in T_a \mathcal{M}, \quad u = \sum_{\mu=1}^m u^\mu U_\mu.$$

Here, U_1, \dots, U_m is a basis of $T_a \mathcal{M}$ so that the components u^μ are exponential coordinates centered at a . Taking partial derivatives with respect to u^μ , it is possible to derive the expressions

$$(2.10) \quad \left. \frac{\partial K_{\delta t, \zeta}}{\partial u^\mu} \right|_{u=0} = 0 \quad \text{and} \quad \left. \frac{\partial^2 K_{\delta t, \zeta}}{\partial u^\mu \partial u^\nu} \right|_{u=0} = -\frac{2}{\|\xi\|^2} \left(g_{\mu\nu} - \zeta \frac{\xi_\mu^* \xi_\nu^*}{\|\xi\|^2} \right).$$

In (2.10), $g_{\mu\nu}$ are the components of the ambient-space induced metric g in (2.3) evaluated in the dual basis U^{*1}, \dots, U^{*m} with $U^{*\mu}(U_\nu) = \delta^\mu_\nu$ [see

also (A.1)]. Moreover, ξ_μ^* are the components of the dual vectors $\xi^* \in T_a^* \mathcal{M}$ given by pulling back the dual vectors to ξ with respect to the canonical inner product (\cdot, \cdot) of \mathbb{R}^n . That is,

$$(2.11) \quad \xi^* := DF^* \Xi = \sum_{\mu=1}^m \xi_\mu^* U^{*\mu},$$

where $DF^* : T_{F(a)}^* \mapsto T_a^* \mathcal{M}$ is the pullback map for dual vectors, and $\Xi = (\xi, \cdot)$. Details of this calculation are provided in Appendix A.2.

Equation (2.10) in conjunction with Lemma 2.1 leads to an asymptotic expression in the sampling interval δt connecting the Hessian of cone kernels to the dual of the dynamical vector field $v^* = g(v, \cdot) = \sum_{\mu=1}^m v_\mu^* U^{*\mu}$:

$$(2.12) \quad \left. \frac{\partial^2 K_{\delta t, \zeta}}{\partial u^\mu \partial u^\nu} \right|_{u=0} = -\frac{2}{\|v\|_g^2 \delta t^2} \left(g_{\mu\nu} - \zeta \frac{v_\mu^* v_\nu^*}{\|v\|_g^2} \right) + O(\delta t^{p-2}) \quad \text{with} \quad v_\mu^* = \sum_{\nu=1}^m g_{\mu\nu} v^\nu.$$

The metric tensor induced on the data manifold by the cone kernel is in fact proportional to the Hessian [9]; in this case $h_{\mu\nu}$ is given by the negative of the $O(\delta t^{-2})$ coefficient in (2.12). Below, we study the geometry induced on the data by $K_{\delta t, \zeta}$ through the associated diffusion operator.

Remark 2.2. Equations (2.10) and (2.12) are unaltered if one replaces the geometric means involving $\|\xi_i\| \|\xi_j\|$ and $\cos^2 \theta_i$ and $\cos^2 \theta_j$ in (2.9) with the corresponding arithmetic and harmonic means, i.e.,

$$(2.13) \quad K_{\delta t, \zeta}(a_i, a_j) = \exp \left[-\frac{\|\omega_{ij}\|^2}{4} \left(\frac{1}{\|\xi_i\|^2} + \frac{1}{\|\xi_j\|^2} \right) (2 - \zeta \cos^2 \theta_i - \zeta \cos^2 \theta_j) \right].$$

For finite δt , the behavior of the two kernels will generally differ. In particular, (2.9) is large if either $\|\xi_i\|$ or $\|\xi_j\|$ are small, whereas (2.13) is large if both $\|\xi_i\|$ and $\|\xi_j\|$ are small (similarly for the angular terms). Thus, the kernel in (2.13) may have higher discriminating power than (2.9), but at the same time may require a larger number of samples for stable behavior.

Remark 2.3. In practical applications, it may be desirable to introduce an additional scaling parameter analogous to ϵ in the isotropic Gaussian kernel in (2.8), i.e.,

$$(2.14) \quad K_{\delta t, \epsilon, \zeta}(a_i, a_j) = \exp \left(-\frac{\|\omega_{ij}\|^2}{\epsilon \|\xi_i\| \|\xi_j\|} [(1 - \zeta \cos^2 \theta_i)(1 - \zeta \cos^2 \theta_j)]^{1/2} \right).$$

Apart from an unimportant scaling factor in the Hessian, the presence of ϵ in (2.14) does not influence the $\delta t \rightarrow 0$ asymptotics, but provides additional freedom to tune the kernel in situations where one does not have control of the sampling interval.

3. The associated diffusion operator and induced metric tensor.

The classical procedure to construct a diffusion operator for geometric analysis of data from a kernel (e.g., [5, 15, 6, 9]) begins with the introduction of an integral operator acting on scalar functions on the data manifold, which we denote here by $H_{\delta t, \zeta}$ to make explicit the two parameters appearing in cone kernels. Specifically,

$$(3.1) \quad H_{\delta t, \zeta} f(a) = \frac{1}{\delta t^m} \int_{\mathcal{M}} K_{\delta t, \zeta}(a, \cdot) f \mu,$$

where f is a sufficiently smooth scalar function on \mathcal{M} , and μ the volume form of the induced metric g in (2.3). Scaling $H_{\delta t, \zeta} f$ by the normalization factor

$$\rho_{\delta t, \zeta}(a) = H_{\delta t, \zeta} \mathbb{1}(a) = \frac{1}{\delta t^m} \int_{\mathcal{M}} K_{\delta t, \zeta}(a, \cdot) \mu,$$

we obtain the integral operator $\mathcal{P}_{\delta t, \zeta}(a) = H_{\delta t, \zeta} f(a) / \rho_{\delta t, \zeta}(a)$. This operator preserves constant functions, i.e., it is an averaging operator.

Next, let

$$(3.2) \quad \mathcal{L}_{\delta t, \zeta} = (\text{Id} - \mathcal{P}_{\delta t, \zeta}) / \delta t^2.$$

This positive-semidefinite operator can be thought of the generator of $P_{\delta t, \zeta}$. An important property of $\mathcal{L}_{\delta t, \zeta}$ is that it annihilates constant functions, i.e., $\mathcal{L}_{\delta t, \zeta} \mathbb{1} = 0$. Following [23], we refer to such operators as diffusion operators.

For suitably-defined kernels, $\mathcal{L}_{\delta t, \zeta}$ converges to a second-order self-adjoint operator $\Delta_{\zeta} = \lim_{\delta t \rightarrow 0} \mathcal{L}_{\delta t, \zeta}$. This operator induces a geometry on the dataset in the sense that it corresponds to a unique codifferential operator δ_{ζ} and Riemannian metric h such that

$$(3.3) \quad \Delta_{\zeta} = \delta_{\zeta} d, \quad \text{where} \quad (w, df)_h = (\delta_{\zeta} w, f)_h$$

for any smooth 1-form field w and scalar function f . Here, $(\cdot, \cdot)_h$ are the canonical (Hodge) inner products for p -form fields associated with h ; i.e.,

$$(3.4) \quad (w, df)_h := \int_{\mathcal{M}} h^{-1}(w, df) \nu, \quad (\delta_{\zeta} w, f) := \int_{\mathcal{M}} \delta_{\zeta}(w) f \nu,$$

where ν and h^{-1} are the volume form and “inverse metric” associated with h , respectively. In particular, we have the following Lemma:

LEMMA 3.1. *The induced metric tensor h at $a \in \mathcal{M}$ associated with the cone kernels (2.9) with FD accuracy from Lemma 2.1 $p \geq 4$ is given by*

$$(3.5) \quad h = \frac{1}{\|v\|_g^2} \left(g - \zeta \frac{v^* \otimes v^*}{\|v\|_g^2} \right),$$

where g is the ambient-space induced metric in (2.3), and $v^* = g(v, \cdot)$ the dual dynamical vector field with respect to g . Moreover, the volume forms ν and μ of h and g , respectively, are related through the expression

$$(3.6) \quad \nu = (1 - \zeta)^{1/2} \bar{\nu}, \quad \text{with} \quad \bar{\nu} = \mu / \|v\|_g^m.$$

A proof of this Lemma can be found in Appendix A.3.

A corollary of Lemma 3.1 is that cone kernels assign a unique metric to equivalence classes of datasets related by conformal transformations. In particular, we say that the datasets associated with the embeddings $F : \mathcal{M} \mapsto \mathbb{R}^n$ and $\tilde{F} : \mathcal{M} \mapsto \mathbb{R}^{\tilde{n}}$ are conformally equivalent if there exists a positive function $r : \mathcal{M} \mapsto \mathbb{R}_+$ such that the induced metric \tilde{g} associated with \tilde{F} is given by $\tilde{g}|_a = r(a)g|_a$. Because conformally equivalent datasets have the properties $\|v\|_{\tilde{g}}^2 = r\|v\|_g^2$ and $\tilde{v}^* := \tilde{g}(v, \cdot) = rv^*$, it follows from (3.5) that both g and \tilde{g} lead to the same h metric; an assertion made in section 2.1.

Remark 3.1. The volume form ν of h is invariant under all transformations that preserve the ratio in the right-hand side of (3.6). This set of transformations includes conformal transformations as a special case, but also admits more general transformations where changes in Riemannian volume are appropriately compensated by changes in the norm of v . We will return to this point in section 3.1.

A further consequence of (3.5) is that h leads to a contraction of distances in neighborhoods of the data manifold where $\|v\|_g$ is uncharacteristically small. Such regions correspond to metastable dynamical regimes separated by rapid transitions with large $\|v\|_g$. As remarked in section 2.1, the ability of locally-scaled kernels (i.e., cone kernels with $\zeta = 0$) to discriminate between regimes of this type has been found to be highly beneficial in Galerkin reduced dynamical models with chaotic metastability [28].

For our purposes, however, of particular interest is the behavior of h and the associated codifferential and diffusion operators in (3.3) in the limit $\zeta \rightarrow 1$, where the directional influence of the dynamical vector field is maximal. In that limit $h(u, v)$ vanishes for all tangent vectors $u \in T_a\mathcal{M}$, i.e., the induced metric becomes degenerate. Equivalently, the length of the integral curves of v measured with respect to h becomes arbitrarily small. In consequence, the diffusion generated by Δ_1 takes place along the integral curves of the dynamical vector field, as we now discuss.

3.1. Behavior in the $\zeta \rightarrow 1$ limit. Recall that a diffusion operator $\Delta = \delta d$ acting on C^2 scalar-valued functions on an m -dimensional manifold \mathcal{M} is said to be along a vector field $v \in T\mathcal{M}$ if the codifferential δw vanishes for all C^1 1-form fields w lying in the $(m-1)$ -dimensional subspace $S_v \subset C^1 T^*\mathcal{M}$ with $w(v) = 0$ [23]. Intuitively, one thinks of the diffusion process generated by Δ to take place along the integral curves of v . A key property of the diffusion operator Δ_ζ in (3.3) associated with cone kernels is that it is along the dynamical vector field v asymptotically as $\zeta \rightarrow 1$, in the sense of the following Lemma.

LEMMA 3.2. *The codifferential operator δ_ζ associated with cone kernels admits the asymptotic expansion*

$$(3.7) \quad \delta_\zeta w = \frac{1}{1-\zeta} \bar{\delta} w + O((1-\zeta)^0), \quad \bar{\delta} w = -\operatorname{div}_{\bar{\nu}}[w(v)v],$$

where $\text{div}_{\bar{\nu}}$ is the divergence operator associated with the volume form $\bar{\nu}$ in (3.5).

This Lemma is proved in Appendix A.4. An asymptotic expansion for the corresponding diffusion operator in (3.3) follows by setting $w = df$ in (3.7) for some C^2 scalar function f , i.e.,

$$(3.8) \quad \Delta_{\zeta} f = \frac{1}{1-\zeta} \bar{\Delta} f + O((1-\zeta)^0), \quad \bar{\Delta} f = -\text{div}_{\bar{\nu}}[v(f)v].$$

A consequence of Lemma 3.2 is that $\delta_{\zeta} w = O((1-\zeta)^0)$ if $w \in S_v$, but $\delta_{\zeta} w' = O((1-\zeta)^{-1})$ if w' lies outside of S_v . As a result, the norm ratio

$$(3.9) \quad \frac{\|\delta_{\zeta} w\|_h}{\|\delta_{\zeta} w'\|_h} = \frac{(\delta_{\zeta} w, \delta_{\zeta} w)^{1/2}}{(\delta_{\zeta} w', \delta_{\zeta} w')^{1/2}} = O((1-\zeta))$$

tends to zero as $\zeta \rightarrow 1$ for all nonzero square-integrable 1-form fields $w' \notin S_v$ for which $\bar{\delta} w'$ is nonzero. We interpret (3.9) as an asymptotic along- v property of Δ_{ζ} .

Remark 3.2. In general, $\Delta_{\zeta} f = \delta_{\zeta} df = -\text{div}_{\nu} \text{grad}_h f$ depends on the metric through the volume form ν , as well as explicitly through the gradient $\text{grad}_h f = h^{-1}(df, \cdot)$. In the $\zeta \rightarrow 1$ limit, the latter is replaced by the directional derivative $v(f)$ in (3.8), which is metric-independent. Thus, the only metric dependence of the limit operator $\bar{\Delta}$ is through the volume form $\bar{\nu}$. According to Remark 3.1, the latter is invariant under a set of transformations of the data which includes conformal transformations as a subset.

3.2. Discrete formulation. Discrete analogs of the diffusion operator (3.3) arise naturally in the framework of discrete exterior calculus (DEC; e.g., [19, 31, 55, 30]). In this setting, the spaces of scalar functions and 1-form fields appearing in (3.3) are replaced by functions $f(a_i) = f_i$ and $w([a_i a_j]) = w_{ij}$ defined on the vertices a_i and edges $[a_i a_j]$, respectively, of a graph formed by the s sampled states in (2.1). These function spaces are equipped with weighted inner products,

$$(3.10) \quad (f, f')_P = \sum_{i=1}^s \pi_i f_i f'_i, \quad (w, w')_P = \sum_{i,j=1}^s \pi_i P_{ij} w_{ij} w'_{ij} / 2,$$

which are the discrete counterparts of (3.4). Also, a difference operator \hat{d} is introduced mapping vertex to edge functions via $\hat{d}f([a_i a_j]) = f_j - f_i$. The associated discrete codifferential $\hat{\delta}_{\delta t, \zeta}$ and diffusion operator $L_{\delta t, \zeta}$ (which in the context of cone kernels depend on both δt and ζ) are then defined in direct analogy with (3.3):

$$(3.11) \quad L_{\delta t, \zeta} = \hat{\delta}_{\delta t, \zeta} \hat{d}, \quad (w, \hat{d}f)_P = (\hat{\delta}_{\delta t, \zeta} w, f)_P.$$

Even though the inner products and associated diffusion operator in (3.11) exist independently of a continuous theory, here we seek to construct $L_{\delta t, \zeta}$ such

that, asymptotically as $s \rightarrow \infty$, it inherits the conformal invariance and along- v properties established in Lemmas 3.1 and 3.2. To that end, we employ the diffusion map (DM) algorithm of Coifman and Lafon [15]. In DM, the inner product weights P_{ij} in (3.10) are the elements of a Markov matrix whose state space is the discrete dataset in (2.1), constructed via the sequence of operations (3.12)

$$K_{ij} = K_{\delta t, \zeta}(a_i, a_j), \quad \tilde{K}_{ij} = \frac{K_{ij}}{(\sum_{k=1}^s K_{ik})^\alpha (\sum_{k=1}^s K_{jk})^\alpha}, \quad P_{ij} = \frac{\tilde{K}_{ij}}{D_{ii} D_{jj}}.$$

Here, α is a real parameter, and D a diagonal degree matrix with $D_{ii} = \sum_{k=1}^s \tilde{K}_{ik}$. The inner-product weights π_i are given by the invariant distribution of that Markov matrix, i.e., $\sum_{i=1}^s \pi_i P_{ij} = \pi_j$. With these definitions, it follows that

$$(3.13) \quad \hat{\delta}_{\delta t, \zeta} w(a_i) = \sum_{j=1}^s P_{ij} (w([a_j a_i]) - w([a_i a_j]))/2, \quad L_{\delta t, \zeta} = I - P.$$

The normalization step to obtain \tilde{K}_{ij} in (3.12), which is not present in standard graph Laplacian algorithms (e.g., [6]), controls the influence of the sampling density with respect to the volume form ν in approximations of the manifold integral in (3.1) by discrete sums of the form $\sum_{j=1}^s \tilde{K}_{ij} f(a_j)$. In particular, under relatively weak assumptions on \mathcal{M} and the embedding map F , it can be shown [15] that for $\alpha = 1$, $\delta t \rightarrow 0$, and $s^{1/2} \delta t^{2+m/2} \rightarrow \infty$ [49], the discrete diffusion operator converges pointwise to Δ_ζ , in the sense that $|\Delta_\zeta f(a) - L_{\delta t, \zeta} f(a)/\delta t^2| \rightarrow 0$ for all $a \in \mathcal{M}$ and sufficiently smooth f . This feature is particularly desirable for our purposes, for the sampling density is dictated by the dynamical flow Φ_t and observation function F , and we are interested in targeting $\mathcal{L}_{\delta t, \zeta}$ without using a priori information about Φ_t and F . We therefore adopt $\alpha = 1$ DM in all of the experiments of sections 4 and 5 ahead.

3.3. Diffusion eigenfunctions. As mentioned in the introduction, diffusion operators such as $L_{\delta t, \zeta}$ in (3.13) are useful for a wide range of data analysis tasks, including dimension reduction, feature extraction, and regularization [48, 5, 14, 15, 6, 33, 7, 42, 13, 16]. Below, we focus on a particular aspect of $L_{\delta t, \zeta}$, namely its eigenfunctions ϕ_i , defined through

$$(3.14) \quad L_{\delta t, \zeta} \phi_i = \lambda_i \phi_i, \quad \phi_i = (\phi_{1i}, \dots, \phi_{si})^T, \quad 0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$$

Diffusion eigenfunctions are traditionally used to create low-dimensional parameterizations of data of the form $a \in \mathcal{M} \mapsto \phi(a) = (\phi_{i_1}(a), \dots, \phi_{i_l}(a)) \in \mathbb{R}^l$, with rigorous embedding results established in the continuous limit [8, 33, 45]. A somewhat different perspective, adopted in NLSA algorithms [25, 28, 29], is to associate low-dimensional subspaces spanned by the leading ϕ_i with spaces of temporal patterns through the time series

$$(3.15) \quad \tilde{\phi}_i(t_j) = \phi_i(a_j) = \phi_{ji},$$

and (in the spirit of [3]) extract spatiotemporal modes of variability through singular value decomposition of the data projected onto those eigenfunctions. Clearly, in both approaches, the relationship of the basis to the underlying dynamical flow plays a major role on algorithm performance.

To gain insight on the influence of the along- v property of diffusion operators associated with cone kernels on their eigenfunctions, it is useful to consider the Dirichlet form associated with Δ_ζ ,

$$\mathcal{E}_\zeta(f) = (f, \Delta_\zeta f)_h.$$

In particular, it follows from the asymptotic expansion in (3.7), in conjunction with the divergence theorem and the assumption in section 2 that \mathcal{M} has no boundary, that

$$(3.16) \quad \mathcal{E}_\zeta(f) = (1 - \zeta)^{-1/2} \bar{\mathcal{E}}(f) + O((1 - \zeta)^{3/2}), \quad \text{with} \quad \bar{\mathcal{E}}(f) = \int_{\mathcal{M}} [v(f)]^2 \bar{\nu}.$$

Therefore, as $\zeta \rightarrow 1$, \mathcal{E}_ζ assigns low energy to functions which (on average) have large directional derivative $v(f)$ along the dynamical flow. Because $\mathcal{E}_\zeta(\phi_i)$ is equal to the corresponding eigenvalue λ_i for normalized eigenfunctions with $\|\phi_i\|_h = 1$, we expect the leading (small- λ_i) diffusion eigenfunctions to vary predominantly in directions transverse to the integral curves of v . This property implies strong adaptivity of ϕ_i to the dynamical flow, and is also independent of the embedding map F and the associated induced metric in (2.3). Below, we demonstrate these properties in numerical experiments.

4. Dynamical systems on the 2-torus. To explicitly illustrate the key features of cone kernels, we begin with a low-dimensional application where the phase space manifold \mathcal{M} is the 2-torus. Denoting by (θ^1, θ^2) the azimuthal and polar angles on the 2-torus, respectively, we consider a two-parameter family of dynamical vector fields

$$(4.1) \quad v = \sum_{\mu=1}^2 v^\mu \frac{\partial}{\partial \theta^\mu}, \quad \text{with} \quad v^1 = 1 + (1 - \beta)^{1/2} \cos \theta^1, \quad v^2 = \Omega(1 - (1 - \beta)^{1/2} \sin \theta^2).$$

Here, Ω is a positive frequency parameter which is set to an irrational number to produce a dense cover of the torus. Moreover, the parameter $\beta \in (0, 1]$ controls the nonlinearity of the flow. Specifically, $\beta = 1$ corresponds to a linear flow, but when $\beta < 1$ the flow “slows down” at $(\theta^1, \theta^2) \sim (\pi, \pi/2)$ and “speeds up” at $(\theta^1, \theta^2) \sim (0, -\pi/2)$. The orbit $(\dot{\theta}^1, \dot{\theta}^2) = (v^1, v^2)$ passing through $(\theta^1, \theta^2) = (0, 0)$ at time $t = 0$ is given by

$$\tan(\theta^1/2) = [1 + (1 - \beta)^{1/2}] \beta^{-1/2} \tan(\beta t/2), \quad \cot(\theta^2/2) = (1 - \beta)^{1/2} + \beta^{1/2} \cot(\beta^{1/2} t/2).$$

Setting $\beta = 0.5$ throughout, we consider the cases $\Omega = 30^{1/2}$ and $30^{-1/2}$, referred to here as Models I and II, respectively. The resulting trajectories in




FIG. 2. Sample trajectories for the dynamical systems (4.1) on 2-tori. (a) Model I: $\Omega = 30^{1/2}$, no deformation. (b) Model II: $\Omega = 30^{-1/2}$, no deformation. (c) Model I': $\Omega = 30^{1/2}$, non-conformal deformation with $\gamma = 0.3$.

data space corresponding to the standard embedding of the 2-torus,

$$F : \mathcal{M} \mapsto \mathbb{R}^3, \quad F(a) = (x^1, x^2, x^3),$$

$$x^1 = (1 + R \cos \theta^2(a)) \cos \theta^1(a), \quad x^2 = (1 + R \cos \theta^2(a)) \sin \theta^1(a), \quad x^3 = \sin \theta^2(a),$$

with $R = 1/2$ are illustrated in Figures 2(a) and 2(b), respectively. In addition to the standard embedding, we also study a non-conformally deformed embedding $F_\gamma(a) = (y^1, y^2, y^3)$ of the 2-torus into \mathbb{R}^3 , where the y^3 coordinate is a stretched version of x^3 with a non-uniform scaling factor, i.e.,

$$y^1 = x^1, \quad \tilde{y}^2 = x^2, \quad y^3 = x^3 e^{\gamma z}, \quad z = (1 + R - x^1)(1 + x^3).$$

The deformed system with $\Omega = 30^{1/2}$ and $\gamma = 0.3$, which we refer to as Model I', is shown in Figure 2(c). In all cases, we equip the ambient data space with the canonical Euclidean inner product.

For each of the models in Figure 2, we generated datasets with $s = 64,000$ samples taken at a sampling interval $\delta t = 2\pi/(S \min\{1, \Omega\})$ where $S = 500$ is an integer parameter controlling the number of samples in each quasi-period. We computed the FD approximations of v using the central scheme in (2.7) with fourth-order accuracy as required by Lemma 2.1 [i.e., the FD weights are $(w_{-2}, \dots, w_2) = (1/12, -2/3, 0, 2/3, -1/12)$]. With this FD scheme, the phase space velocity norm ratio $\max\{\|\xi_i\|\}/\min\{\|\xi_i\|\}$ for Models I, II, and I' was of order 20, 35, and 40, respectively. Thus, the influence of the $\|\xi_i\|/\|\xi_j\|$ scaling factors in the cone kernels (2.9) is expected to be significant.

We evaluated the discrete diffusion operator $L_{\delta t, \zeta}$ from cone kernels with $\zeta = 0$ and 0.995 (corresponding to no influence and strong influence of the directionality of ξ , respectively) using the DM procedure in (3.12) with $\alpha = 1$. For comparison, we also computed the diffusion operator associated with the isotropic Gaussian kernel in (2.8) with $\epsilon = 0.1$. To limit memory usage, in all cases we truncated the pairwise kernel evaluations to $b = 2000$ nearest neighbors [in the sense of K_{ij} in (3.12)] per sample. The kernel values at truncation

were no greater than $O(10^{-7})$, indicating that truncation has negligible impact on the numerical results.

Representative eigenfunctions for Models I, II, and I' are shown in Figures 3–5. To test for convergence of our results to the continuous limit, we performed a series of long runs, where s was increased eightfold, δt decreased twofold, and b set to 10,000 (in the isotropic Gaussian kernel case ϵ was decreased fourfold). The diffusion eigenfunctions from the long runs (not shown here) were in good agreement with those displayed in Figures 3–5. We also performed eigenfunction calculations using a first-order backward FD scheme so that $\xi_i = X_i - X_{i-1}$, and found again only minor changes relative to the results with fourth-order FD accuracy. It therefore appears that in this setting with dense sampling the accuracy required by Lemma 3.1 is not crucial to recover the salient features of the eigenfunctions.

4.1. Properties of the diffusion eigenfunctions. First, consider the eigenfunctions for Models I and II. According to [15], the results of DM with the isotropic Gaussian kernel should only depend on the embedding F , and should be independent of the sampling density on \mathcal{M} induced by the dynamical flow. Indeed, as expected from theory, the isotropic-kernel eigenfunctions in Figures 3 and 4 are essentially equivalent, despite the fact that the underlying dynamical system trajectories are qualitatively different.

On the other hand, because cone kernels depend explicitly on the dynamical flow through ξ [see (2.9)], the eigenfunctions for Models I and II differ. In the $\zeta = 0$ results for Model I shown in Figure 3, the eigenfunction wavecrests are compressed in the portion of the torus where the flow evolves slowly, and rarefied in the transitory region characterized by large $\|v\|_g$ (compare, e.g., the ϕ_3 and ϕ_5 eigenfunctions obtained via the isotropic kernel and the $\zeta = 0$ cone kernel). As a result, these eigenfunctions have higher discriminating power in the regions of \mathcal{M} where the system evolves more slowly. The Model II eigenfunctions exhibit a qualitatively similar behavior.

The along- v property of the diffusion operators with $\zeta \approx 1$, and the resulting adaptation of the eigenfunctions to the dynamical flow expected on the basis of (3.16), is manifested in the right-hand columns of Figures 3 and 4. There, the eigenfunctions vary predominantly in directions transverse to the dynamical flow, resulting in characteristic swirl patterns for Model I and azimuthal streak-like patterns for Model II.

Because the integral curves of v are metric independent, this property is robust against changes in the embedding function F . Indeed, in the Model I' results in Figure 5, it is only the $\zeta = 0.995$ eigenfunctions which remain qualitatively similar to those in Figure 3. There, the dataset deformation has left a clear imprint on the eigenfunctions. On the other hand, the $\zeta = 0.995$ eigenfunctions retain the dynamics-adaptation featuring the characteristic swirl patterns following the dynamical flow. The structure of the eigenfunctions transversely to the flow does exhibit some changes relative to Figure 3, but

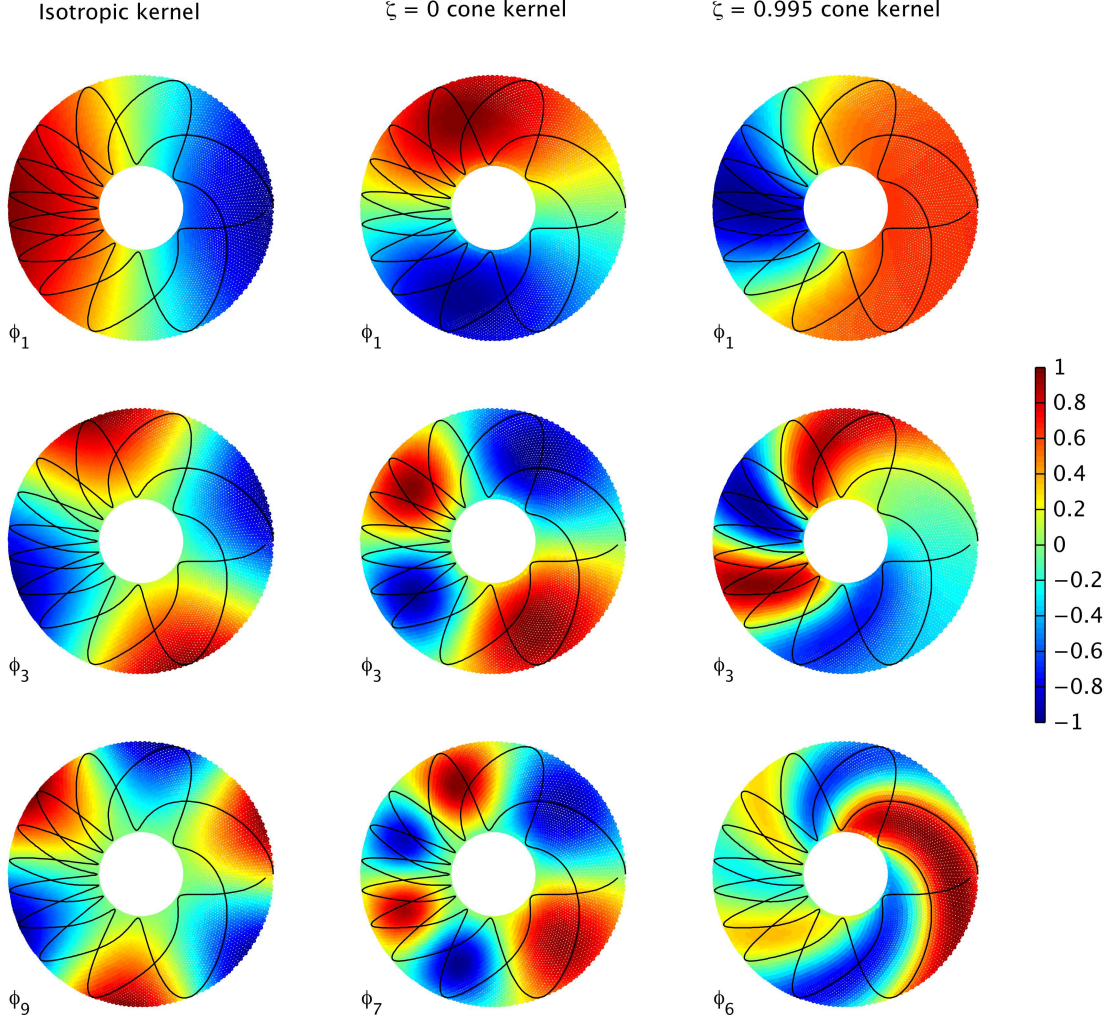


FIG. 3. Scatterplots of diffusion eigenfunctions for the dynamical system on the 2-torus with $\Omega = 30^{1/2}$ obtained using DM with $\alpha = 1$ in conjunction with the isotropic Gaussian kernel (2.8), and the cone kernels (2.9) with $\zeta = 0$ and 0.995. A portion of the dynamical system trajectory is plotted in a black line for reference.

these are significantly weaker compared to the isotropic Gaussian kernel and $\zeta = 0$ cone kernel examples.

5. North Pacific SST data from a comprehensive climate model.

In this experiment, we apply cone kernels to SST data in the North Pacific sector of the Community Climate System Model version 3 [17]. The dataset, which was studied in [26] via NLSA, consists of monthly-averaged samples of SST in the rectangular domain 120°E – 110°W and 20°N – 65°N spanning a

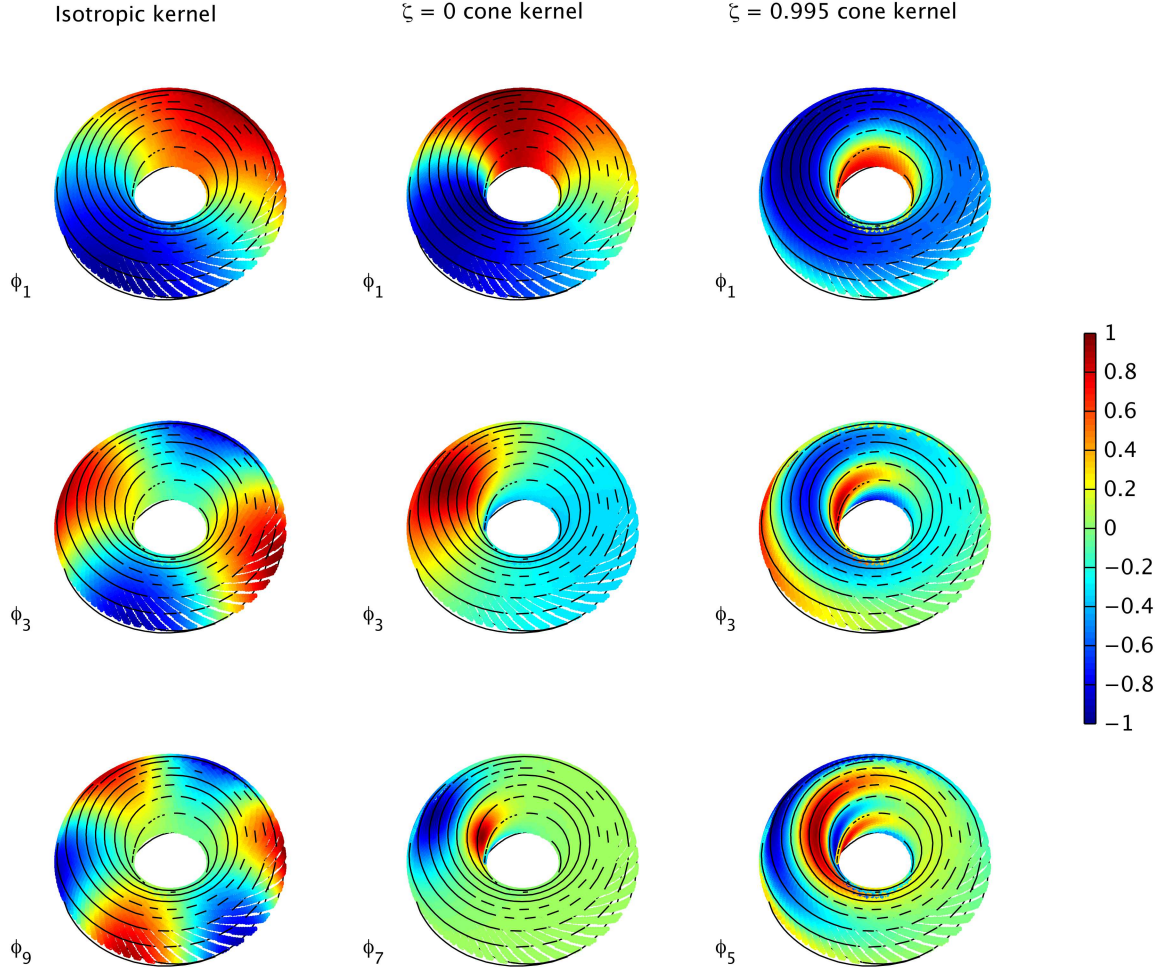


FIG. 4. Same as Figure 3 but for the dynamical system with $\Omega = 30^{-1/2}$.

900-year interval.¹ We work throughout with the model's nonuniform native ocean grid of 1° nominal horizontal resolution. The number of gridpoints in the analysis domain is $d = 6671$.

A major component of North Pacific SST variability is due to the annually varying solar forcing (the seasonal cycle). The latter is superposed to low-frequency (interannual to decadal) variability patterns, the most prominent of which are the Pacific Decadal Oscillation (PDO) [41] and the North Pacific Gyre Oscillation (NPGO) [21]. A signature of ENSO (which is most prominent in the tropical Pacific [53]) is also present in this domain. We refer the reader to [26] for further details on these modes of variability extracted via NLSA. Here, our objective is to compare the diffusion eigenfunctions of cone kernels

¹The dataset is available at the Earth System Grid repository, <http://www.earthsystemgrid.org>, where it is designated CCSM3 integration b30.004.




FIG. 5. Same as Figure 3 but for the deformed torus with deformation parameter $\gamma = 0.3$ [see Figure 2(c)].

with $\zeta \approx 1$ to those with $\zeta = 0$, which are equivalent to the earlier NLSA kernels (aside from the fourth-order accurate FD scheme used here in place of the first-order backward scheme in [26]).

Following [26, 29], to “Markovianize” the time-dependent solar forcing in the data, and induce timescale separation in the diffusion eigenfunctions, we first embed the data to a higher-dimensional space via Takens’ method of delays [44, 51, 11, 47]. That is, we map each spatial snapshot $x_i \in \mathbb{R}^d$ to a spatiotemporal sequence $X_i = (x_i, x_{i-1}, \dots, x_{i-(q-1)}) \in \mathbb{R}^n$, where q is an integer parameter measuring the embedding window length (in months). Here, we work with a two-year embedding window, $q = 24$, which corresponds to a data space dimension $n = qd = 160,104$. Qualitatively similar




FIG. 6. Eigenvalues λ_i for CCSM North Pacific SST data obtained via cone kernels with $\zeta = 0$ and 0.995, normalized so that $1 - \lambda_1 = 1$. To highlight the differences between the two spectra, $1 - \lambda_i$ values are plotted in a logarithmic scale. Periodic, low-frequency, and intermittent eigenfunctions are indicated using \circ , \triangle , and \square markers, respectively. Solid markers correspond to eigenfunctions whose temporal character does not belong in these families.

results can be obtained with embedding windows in the interval 1–5 years. Similarly to the torus examples of section 4, we evaluate ξ using a central fourth-order FD scheme. The number of samples s available for analysis after embedding and removal of two samples in the beginning and end of the simulation interval (to compute ξ_1 and ξ_s) is 10,773. Unlike the torus experiments, s is small-enough in this case so as not to require nearest-neighbor truncation. We employ an area-weighted data space inner product given by $(X_i, X_j) = \sum_{\nu=1}^n A_{\nu \bmod d} X_i^\nu X_j^\nu$, where A_1, \dots, A_d are the grid cell areas in the analysis domain. Prior to analysis we center the data by subtracting the temporal mean $\bar{x} = \sum_i x_i$ from each snapshot x_i to produce temperature anomaly snapshots. Note that centering is done mainly for visualization purposes, and does not influence the kernel values in (2.8) and (2.9).

Representative diffusion eigenvalues λ_i and eigenfunctions ϕ_i obtained via cone kernels with $\zeta = 0$ and 0.995 are displayed in Figures 6–8, where the ϕ_i are represented by the corresponding $\tilde{\phi}_i$ time series from (3.15). To visualize subspaces spanned by selected groups of eigenfunctions $\phi = (\phi_{i_1}, \dots, \phi_{i_l})$ as spatiotemporal patterns, we filter the data using the $s \times l$ matrix ϕ as a convolution filter; i.e., $X \mapsto X\pi\phi\phi^T$, where $\pi = \text{diag}(\pi_1, \dots, \pi_s)$ is the diagonal matrix formed by the inner product weights in (3.10). Spatiotemporal patterns of this type are shown in Figure 9 and Movie 1, which is much more revealing. Note that the eigenfunctions computed using a first-order backward FD scheme (not shown here) exhibit minor changes compared to those in Figures 7 and 8.

Remark 5.1. In this application, the variations in the phase space ve-




FIG. 7. Representative eigenfunctions for CCSM North Pacific SST data obtained via the cone kernel with $\zeta = 0$. The left-hand panels show 20-year portions of the time series $\tilde{\phi}_i(t) = \phi_i(a_t)$. The right-hand panels are frequency spectra computed via the discrete Fourier transform of $\tilde{\phi}_i$. The eigenfunctions in this Figure have been chosen so as to qualitatively match as close as possible those of Figure 8. Note the timescale mixing in $\{\phi_{21}, \phi_{22}\}$ compared to the corresponding $\zeta = 0.995$ eigenfunctions, $\{\phi_{13}, \phi_{14}\}$, in Figure 8.

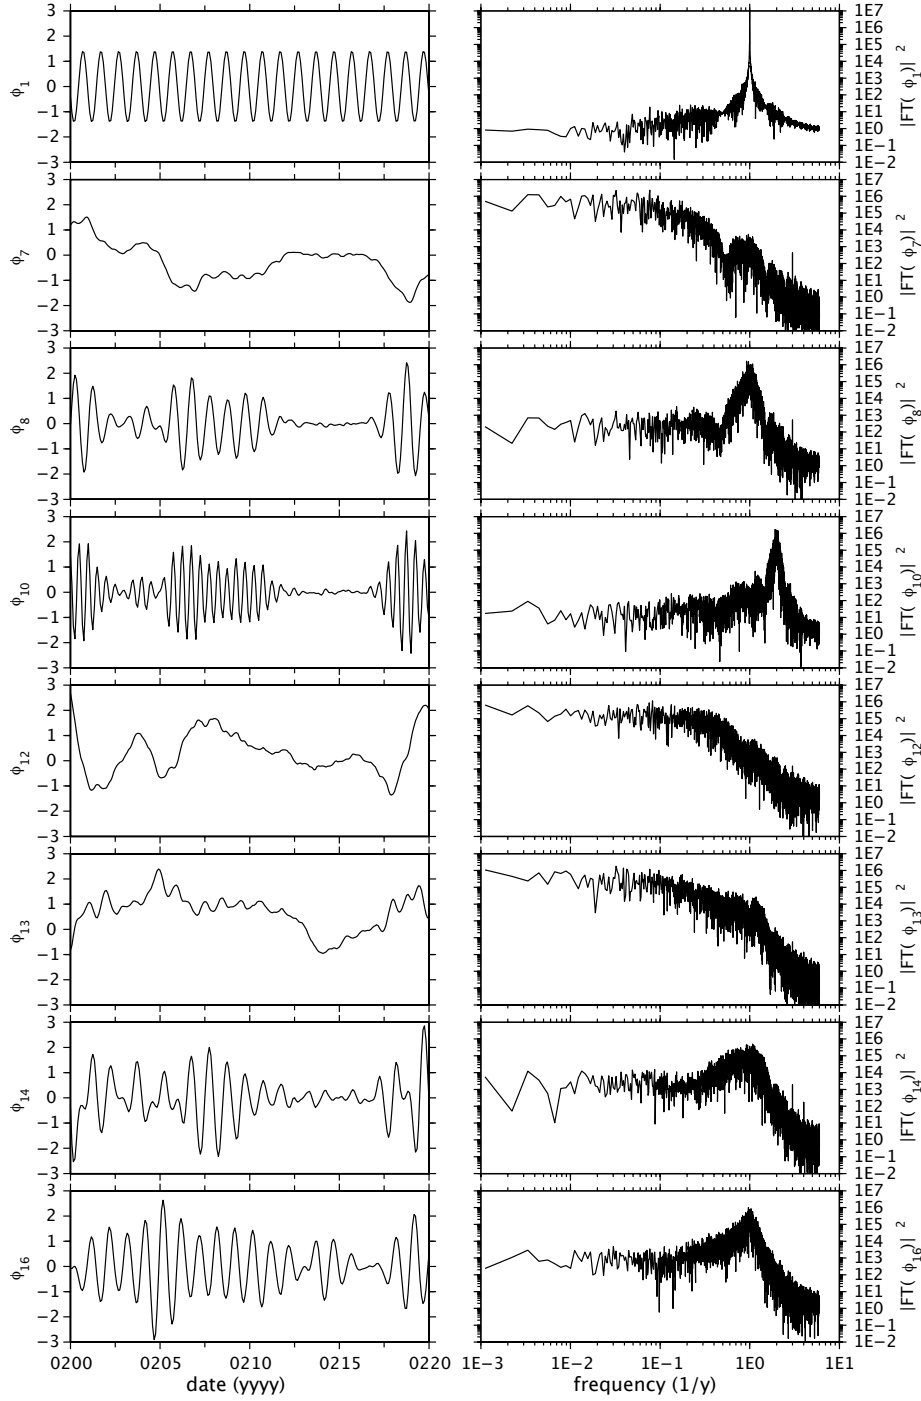


FIG. 8. Same as Figure 7 but for the cone kernel with $\zeta = 0.995$. The displayed eigenfunctions are (a) annual periodic; (b) PDO; (c,d) annual and semiannual intermittent patterns associated with the PDO; (e) ENSO; (f) NPGO; (g) annual intermittent patterns associated with ENSO; (h) annual intermittent patterns associated with the NPGO.




FIG. 9. Snapshots of spatiotemporal patterns of SST anomalies (in K) for September of simulation year 207 obtained from the eigenfunctions in Figures 7 and 8. The displayed patterns are the (a) annual periodic; (b) PDO; (c,d) annual and semiannual intermittent patterns associated with the PDO; (e) ENSO; (f) NPGO; (g) annual intermittent patterns associated with ENSO; (h) annual intermittent patterns associated with the NPGO. Note the higher amplitude in panels (f-h) for the $\zeta = 0.995$ cone kernel compared to the $\zeta = 0$ cases. See Movie 1 for the dynamic evolution of these patterns and the raw data.

locity velocity norm $\|\xi_i\|$ are of order 10% (cf. section 4). As a result, the influence of the local scaling by $\|\xi_i\|\|\xi_j\|$ in this case is rather weak, and correspondingly the eigenfunctions obtained via the $\zeta = 0$ cone kernel and the isotropic Gaussian kernel (not shown here) do not differ significantly. Here, the main difference between the isotropic and cone kernels is due to the directionality of the dynamical flow, which is influenced strongly by the annual cycle.

5.1. Spatiotemporal patterns extracted by diffusion eigenfunctions. In general, the eigenfunctions can be grouped into three families according to the timescales present in the $\tilde{\phi}_i$ time series, which we refer to as periodic, low-frequency, and intermittent [26]. Figures 7 and 8 display examples from each family. The periodic eigenfunctions come in doubly-degenerate pairs, and are dominated by a single frequency component which is an integer multiple of the annual cycle. Due to the presence of strong seasonality, one would expect the data manifold to have the structure of a circle along one of its dimensions—the presence of the periodic eigenfunctions is consistent with this picture. The low-frequency family is characterized by red-noise type frequency spectra, featuring significant power at interannual to decadal timescales. These basis functions describe familiar low-frequency patterns of North Pacific SST variability, including the PDO, ENSO, and NPGO [Panels (b), (e), and (f) in Figure 9 and Movie 1], which are also accessible through SSA algorithms or PCA of seasonally-detrended data. The intermittent eigenfunctions have the structure of amplitude-modulated wavetrains consisting of a periodic carrier signal with integer frequency multiples of the seasonal cycle, which is modulated by a low-frequency envelope. In certain cases (though not always) the low-frequency envelope is correlated strongly with a low-frequency eigenfunction. In Figure 8, for instance, the low-frequency envelopes of the intermittent eigenfunctions $\{\phi_8, \phi_{10}\}$, ϕ_{14} , and ϕ_{16} correlate strongly with the low-frequency eigenfunctions ϕ_7 , ϕ_{12} , and ϕ_{13} , respectively.

Similarly to the periodic eigenfunctions, the intermittent eigenfunctions arise in near-degenerate pairs (see Figure 6). The corresponding spatiotemporal patterns [Panels (c), (d), (g), and (h) in Figure 9 and Movie 1] feature propagating structures in regions of the North Pacific with high variability, including the Kuroshio current, the Bering Sea, and the west coast of North America. These patterns carry approximately two orders of magnitude less variance of the raw data than the prominent low-frequency modes. As a result, they are not accessible to variance-greedy PCA-type algorithms. However, these eigenfunctions are crucial in explaining the lagged temporal correlation structure of the data in so-called anomaly reemergence mechanisms [1, 12], and have also been found to have high skill as external factors in regression models for SST variability [27].

5.2. Influence of the angular term. The periodic, low-frequency, and intermittent families arise in both of the $\zeta = 0$ and 0.995 kernels. As is evident

in Figures 7 and 8, there exists a reasonably clear qualitative correspondence between the leading eigenfunctions in each case. Nevertheless, the two spectra differ in two important aspects.

First, the number of eigenfunctions required to capture the prominent features described in section 5.1 is significantly smaller in the $\zeta = 0.995$ case than the $\zeta = 0$ example. In particular, it takes 17 $\zeta = 0.995$ eigenfunctions to describe the annual and semiannual cycles, the PDO, ENSO, and NPGO low-frequency modes, and the associated intermittent modes, versus 26 $\zeta = 0$ eigenfunctions. The extra eigenfunctions contained in the $\zeta = 0$ spectrum are generally associated with higher harmonics of the seasonal cycle, which appear to be of limited physical significance. Examples of these harmonics can be seen in the eigenvalue spectrum of Figure 6, where the annual and semiannual periodic eigenfunctions (ϕ_1 – ϕ_4) are followed by higher harmonics all the way to the Nyquist limit 5 y^{-1} associated with the one-month sampling interval. Higher-frequency harmonics are also present in the intermittent eigenfunctions.

The suppression of the high-frequency harmonics from the $\zeta = 0.995$ spectrum is consistent with the observation made in section 3.3 that large gradients of the eigenfunctions along the direction of the dynamical vector field v lead to large Dirichlet form in (3.16), and correspondingly large eigenvalues. Because a major component of v is due to the annual cycle, the higher-harmonics incur a large $v(\phi_i)$ penalty, and thus are removed from the leading part of the spectrum. We expect this to be a generic feature in datasets with prominent quasi-periodic behavior.

A further key difference between the $\zeta = 0$ and $\zeta = 0.995$ eigenfunctions pertains to timescale separation. In particular, as can be seen in Figure 7, the ϕ_{22} eigenfunction for $\zeta = 0$ representing the NPGO exhibits a mixture of interannual and annual timescales. Likewise, ϕ_{21} has a mixed intermittent–low-frequency character. On the other hand, the corresponding $\zeta = 0.995$ eigenfunctions in Figure 8 (ϕ_{13} and ϕ_{14} , respectively) have a well-separated temporal character, with ϕ_{13} acting as a low-frequency envelope of ϕ_{16} . In consequence, the corresponding spatiotemporal patterns have larger amplitude and feature more prominent intermittent coherent structures; see, e.g., the portion of the dynamic evolution around simulation year 207 in Figure 9 and Movie 1. Again, we attribute this behavior to the superior adaptation of the $\zeta = 0.995$ eigenfunctions to the dynamical flow.

6. Conclusions. In this work, we have developed a family of kernels for data analysis in dynamical systems, which incorporate empirical information about the dynamical flow in phase space. Compared to canonical isotropic Gaussian kernels, the kernels presented here assign higher affinity to pairs of data samples whose relative displacement vector lies on a small-angle cone with axis parallel to the dynamical vector field v on the data manifold. The latter is estimated through finite differences of time-ordered samples, i.e., without requiring prior knowledge of the equations of motion. Due to the presence of

the angular term, we refer to the new kernels as cone kernels.

Cone kernels also feature a scaling factor introduced heuristically in earlier work on so-called NLSA algorithms [28, 29], whose role is to locally decrease (increase) the rate of decay of the kernel in regions of phase space where the dynamical flow is fast (slow) in the sense of the norm $\|v\|_g$ of the dynamical vector field in data space. Moreover, the strength of the angular dependence is controlled by a parameter $\zeta \in [0, 1)$, such that $\zeta = 0$ and $\zeta \rightarrow 1$ correspond to zero or maximal influence of the angular term. Thus, cone kernels include the earlier NLSA kernels as the special case $\zeta = 0$.

We evaluated the metric tensor h induced on the data manifold by cone kernels (Lemma 3.1) using the asymptotic analysis framework of Berry [9], and also studied the associated diffusion operator Δ_ζ . By virtue of the $\|v\|_g$ -dependent scaling, the induced metric is invariant for all ζ under conformal transformations of the original ambient space metric g . Moreover, for $\zeta > 0$, h contracts local distances between points on the data manifold whose relative displacement is parallel to v , becoming degenerate as ζ approaches 1. In that regime, Δ_ζ becomes along v [23], in the sense that the associated codifferential operator asymptotically annihilates all differential 1-forms w with the property $w(v) = 0$ (Lemma 3.2).

Intuitively, one thinks of Δ_ζ in the $\zeta \rightarrow 1$ limit as generating diffusions along the integral curves of the dynamical vector field. Because v and its integral curves do not depend on the ambient space metric, this feature is intrinsic to the dynamical system under study. More generally, as $\zeta \rightarrow 1$, the action of Δ_ζ on functions depends on the ambient space metric only through the ratio $\mu/\|v\|_g^m$ for an m -dimensional manifold. The latter is invariant under conformal transformations of g , as well as more general transformations. A further important property arising in the $\zeta \rightarrow 1$ limit is that the Dirichlet form associated with Δ_ζ depends on the directional derivative of functions along v , as opposed to the canonical dependence on the gradient operator [see (3.16)]. This property has significant bearing on the structure of the corresponding diffusion eigenfunctions, which are useful in a wide range of dimension reduction, signal processing, and learning problems. In particular, the leading eigenfunctions are expected to be adapted to the dynamical system generating the data, in the sense of varying weakly along the integral curves of v .

We discussed the utility of cone kernels in a suite of numerical experiments involving nonlinear flows on the 2-torus and North Pacific sea surface temperature (SST) data from a comprehensive climate model. In the torus experiments, we explicitly demonstrated the adaptivity of the diffusion eigenfunctions associated with $\zeta \approx 1$ cone kernels to the dynamics (Figures 3 and 4), as well as the robustness of those eigenfunctions to non-conformal deformations of the data (Figure 5). In the North Pacific SST experiments, cone kernels were found to have superior feature extraction capabilities, in the sense of requiring fewer basis functions than their $\zeta = 0$ counterparts to describe the

salient coherent structures of North Pacific SST variability, while also providing better separation between the timescales associated with the annual solar forcing and low-frequency (interannual) variability of the ocean. We attribute this improvement of skill to the ability of $\zeta \approx 1$ cone kernels to take into account changes in the direction of the dynamical flow due to the annual cycle. This feature should be generic in datasets with prominent quasiperiodic behavior.

There are several open questions generated by this work which lie outside the scope of the present paper. On the theory side, a more detailed understanding of the diffusion eigenfunctions obtained from cone kernels would be desirable; e.g., their properties as embedding coordinates [8, 33, 45]. Furthermore, it would be useful to explore generalizations of the deterministic framework adopted here to stochastic dynamical systems. Such approaches might involve replacing the ambient-space distances and inner products with suitable statistical metrics [52], retaining the explicit dependence on aspects related to the drift of the system in phase space. Even in the deterministic dynamical system context, potential shortcomings of cone kernels may arise due to sensitivity to observational noise and/or poor performance of the directional term in high intrinsic dimensions. Such scenarios would warrant modification of the cone kernel formulation put forward here, but we expect the general approach of incorporating empirically accessible information about the dynamics in kernel design to remain fruitful.

Appendix. Technical results. This Appendix contains the proof of Lemmas 2.1, 3.1, and 3.2, as well as derivations of a number of results used in the main text. Whenever convenient, we use the shorthand notation ∂_μ to represent partial differentiation $\partial/\partial u^\mu$ with respect to a manifold coordinate u^μ .

A.1. Proof of Lemma 2.1. Let (u^1, \dots, u^m) be a coordinate system defined in a neighborhood of $a_i \in \mathcal{M}$, and U_1, \dots, U_m be the corresponding coordinate basis vectors of $T_{a_i}\mathcal{M}$. Denote the coordinates of the curve $\Phi_t a_0$ by $u^\mu(t)$, so that

$$v|_{a_i} = \sum_{\mu=1}^m v^\mu U_\mu, \quad \text{with} \quad v^\mu = \left. \frac{du^\mu}{dt} \right|_{t_i}.$$

Moreover, fix a basis e_1, \dots, e_n of \mathbb{R}^n so that $X_i = F(a_i) = \sum_{\nu=1}^n X^\nu e_\nu$. The $\{e_\nu\}$ basis can be identified with a basis of $T_{X_i}\mathbb{R}^n$ via the canonical isomorphism $\mathbb{R}^n \simeq T_{X_i}\mathbb{R}^n$. With this choice of bases, the components of the derivative map $DF : T_{a_i}\mathcal{M} \mapsto T_{X_i}\mathbb{R}^n$ become $DF^\nu{}_\mu = \partial X^\nu / \partial u^\mu|_{a_i}$. We then

compute

$$\begin{aligned}\xi_i &= \delta_p X_i = \left. \frac{dX}{dt} \right|_{t_i} \delta t + O(\delta t^{p+1}) = \left(\sum_{\mu=1}^m \frac{\partial X}{\partial u^\mu} \frac{du^\mu}{dt} \right) \Big|_{t_i} \delta t + O(\delta t^{p+1}) \\ &= \sum_{\mu=1}^m \sum_{\nu=1}^n \frac{\partial X^\nu}{\partial u^\mu} \Big|_{a_i} v^\mu e_\nu \delta t + O(\delta t^{p+1}) = DF v|_{a_i} \delta t + O(\delta t^{p+1}),\end{aligned}$$

which gives (2.5). Equation (2.6) follows immediately:

$$\begin{aligned}\|v\|_{g,a_i}^2 &= (DF v, DF v)|_{a_i} = (\xi_i + O(\delta t^{p+1}), \xi_i + O(\delta t^{p+1}))/\delta t^2 \\ &\implies \|v_i\|_g = \|\xi_i\|/\delta t + O(\delta t^p). \quad \square\end{aligned}$$

A.2. Derivatives of cone kernels. Let $\{e_1, \dots, e_n\}$ be a basis of \mathbb{R}^n and e^{*1}, \dots, e^{*n} its dual basis with $e^{*\mu}(e_\nu) = \delta^\mu_\nu$. Let also $c_{\mu\nu} = (e_\mu, e_\nu)$ be the corresponding matrix elements of the data space inner product, giving the components of Ξ in (2.11) via

$$\Xi = \sum_{\nu=1}^n \Xi_\nu e^{*\nu}, \quad \Xi_\nu = \sum_{\mu=1}^n c_{\mu\nu} \xi^\mu, \quad \text{with} \quad \xi = \sum_{\mu=1}^n \xi^\mu e_\mu.$$

To evaluate the derivatives of cone kernels in (2.10), it is convenient to write down $K_{\delta t, \zeta}(a, \exp_a u) = \exp(-A(u)B(u))$ with

$$\begin{aligned}A(u) &= \frac{1}{\|\xi\|^2} \left(\|\omega\|^2 - \zeta \frac{(\xi, \omega)^2}{\|\xi\|^2} \right), \quad B(u) = \frac{\|\xi\|}{\|\xi'\|} \sqrt{C(u)}, \\ C(u) &= \frac{\|\xi'\|^2 \|\omega\|^2 - \zeta(\xi', \omega)^2}{\|\xi\|^2 \|\omega\|^2 - \zeta(\xi, \omega)^2}, \quad \omega = F(a') - F(a), \quad \xi' = \delta_p F(a').\end{aligned}$$

The function $A(u)$ vanishes at $u = 0$, and its leading two derivatives with respect to u can be computed straightforwardly:

$$\begin{aligned}\frac{\partial A}{\partial u^\mu} &= \frac{2}{\|\xi\|^2} \sum_{\rho, \sigma=1}^n \left(c_{\rho\sigma} \omega^\rho \frac{\partial \omega^\sigma}{\partial u^\mu} - \zeta(\xi, \omega) c_{\rho\sigma} \xi^\rho \frac{\partial \omega^\sigma}{\partial u^\mu} \right) \\ \frac{\partial^2 A}{\partial u^\mu \partial u^\nu} &= \frac{2}{\|\xi\|^2} \sum_{\rho, \sigma=1}^n \left(c_{\rho\sigma} \frac{\partial \omega^\rho}{\partial u^\nu} \frac{\partial \omega^\sigma}{\partial u^\mu} + c_{\rho\sigma} \omega^\rho \frac{\partial^2 \omega^\sigma}{\partial u^\mu \partial u^\nu} \right) \\ &\quad - \frac{\zeta}{\|\xi\|^2} \sum_{\alpha, \beta, \rho, \sigma=1}^n \left(c_{\alpha\beta} c_{\rho\sigma} \xi^\alpha \xi^\beta \frac{\partial \omega^\rho}{\partial u^\nu} \frac{\partial \omega^\sigma}{\partial u^\mu} - (\xi, \omega) c_{\rho\sigma} \xi^\rho \frac{\partial^2 \omega^\sigma}{\partial u^\mu \partial u^\nu} \right).\end{aligned}$$

Noting that $\partial \omega^\nu / \partial u^\mu|_{u=0} = DF^\nu_\mu$, these expressions lead to

$$\left. \frac{\partial A}{\partial u^\mu} \right|_{u=0} = 0, \quad \left. \frac{\partial^2 A}{\partial u^\mu \partial u^\nu} \right|_{u=0} = \frac{2}{\|\xi\|^2} \left(g_{\mu\nu} - \zeta \frac{\xi_\mu^* \xi_\nu^*}{\|\xi\|^2} \right),$$

where

$$(A.1) \quad g_{\mu\nu} = \sum_{\rho, \sigma=1}^n c_{\rho\sigma} DF^\rho{}_\mu DF^\sigma{}_\nu$$

are the components of the ambient spaced induced metric g in (2.3), and ξ_μ^* are defined in (2.11). Therefore, assuming that B and its first two derivatives are all bounded at $u = 0$, we have

$$(A.2) \quad \begin{aligned} \left. \frac{\partial K_{\delta t, \zeta}}{\partial u^\mu} \right|_{u=0} &= -K_{\delta t, \zeta}(a, a) \left(\frac{\partial A}{\partial u^\mu} B + A \frac{\partial B}{\partial u^\mu} \right) \Big|_{u=0} = 0, \\ \left. \frac{\partial^2 K_{\delta t, \zeta}}{\partial u^\mu \partial u^\nu} \right|_{u=0} &= - \left[\frac{\partial K_{\delta t, \zeta}}{\partial u^\mu} \left(\frac{\partial A}{\partial u^\mu} B + A \frac{\partial B}{\partial u^\mu} \right) \right] \Big|_{u=0} \\ &\quad - K_{\delta t, \zeta}(a, a) \left(\frac{\partial^2 A}{\partial u^\mu \partial u^\nu} B + \frac{\partial A}{\partial u^\mu} \frac{\partial B}{\partial u^\nu} + \frac{\partial A}{\partial u^\nu} \frac{\partial B}{\partial u^\mu} + A \frac{\partial^2 B}{\partial u^\mu \partial u^\nu} \right) \Big|_{u=0} \\ &= - \frac{\partial^2 A}{\partial u^\mu \partial u^\nu} B \Big|_{u=0}. \end{aligned}$$

Thus, in order to obtain (2.10), it suffices to determine the value $B(0)$ and check that the first two derivatives of B are bounded at $u = 0$. Now, because ξ is a smooth function of u and $B(0) = \sqrt{C(0)}$, any pathological behavior of B and its derivatives at the origin would be caused by C .

To confirm that this is not the case, consider the non-degenerate type $(0, 2)$ tensor $Q(u)$ with components

$$Q(u)_{\mu\nu} = (\|\xi\|^2 \delta_{\mu\nu} - \zeta \xi_\mu^* \xi_\nu^*)|_u,$$

where ξ_μ^* is given by (2.11). Using $Q(u)$, we expand $C(u)$ as follows:

$$\begin{aligned} C(u) &= \frac{\sum_{\mu, \nu=1}^m Q(u)_{\mu\nu} u^\mu u^\nu}{\sum_{\rho, \sigma=1}^m Q(0)_{\rho\sigma} u^\rho u^\sigma} = 1 + \sum_{\alpha=1}^m u^\alpha R_\alpha^{(1)} + \frac{1}{2} \sum_{\alpha, \beta=1}^m u^\alpha u^\beta R_{\alpha\beta}^{(2)} + O(\|u\|^3), \\ R_\alpha^{(1)} &= \frac{\sum_{\mu, \nu=1}^m u^\mu u^\nu \partial_\alpha Q_{\mu\nu}(0)}{\sum_{\rho, \sigma=1}^m u^\rho u^\sigma Q_{\rho\sigma}(0)}, \quad R_{\alpha\beta}^{(2)} = \frac{\sum_{\mu, \nu=1}^m u^\mu u^\nu \partial_\alpha \partial_\beta Q(0)_{\mu\nu}}{\sum_{\rho, \sigma=1}^m u^\rho u^\sigma Q(0)_{\rho\sigma}}. \end{aligned}$$

The tensor components $R_\alpha^{(1)}$ and $R_{\alpha\beta}^{(2)}$ are bounded in absolute value through the inequalities

$$|R_\alpha^{(1)}| \leq |\Lambda_{\max}^{(1)} / \Lambda_{\min}^{(1)}|, \quad |R_{\alpha\beta}^{(2)}| \leq |\Lambda_{\max}^{(2)} / \Lambda_{\min}^{(2)}|,$$

where $\Lambda_{\max}^{(i)}$ ($\Lambda_{\min}^{(i)}$) are the largest (smallest) generalized eigenvalues

$$\partial_\alpha Q(0)_{\mu\nu} y = \Lambda^{(1)} Q(0)_{\mu\nu} y, \quad \partial_\alpha \partial_\beta Q(0)_{\mu\nu} y = \Lambda^{(2)} Q(0)_{\mu\nu} y$$

in absolute value. Thus, $C(0) = B(0) = 1$, and the leading two derivatives of C and B are all bounded, leading through (A.2) to (2.10).

A.3. Proof of Lemma 3.1. We prove this Lemma following the analysis of Berry [9] connecting the induced metric to the Hessian of the kernel.

First, fix a basis U_1, \dots, U_m of $T_a \mathcal{M}$, and consider the corresponding exponential coordinates u^1, \dots, u^m with respect to the ambient-space induced metric g in (2.3). For small-enough sampling interval δt , the m -dimensional Euclidean ball $B_{\delta t}$ of radius $\|v\|_{g,a} \delta t$ centered at the origin maps diffeomorphically to a neighborhood $\mathcal{B}_{\delta t} = \exp_a B_{\delta t} \subset \mathcal{M}$ of a . In these coordinates, the volume element $d\mu$ associated with g becomes $\mu|_{\exp_a a} = \det D \exp_a|_u dU^{*1} \wedge \dots \wedge dU^{*m}$, where $dU^{*\mu}$ are the dual basis vectors to U_μ , and $D \exp_a|_u : T_u T_a \mathcal{M} \mapsto T_a \mathcal{M}$ is the derivative of the exponential map,

$$D \exp_a|_u(u') = \lim_{\tau \rightarrow 0} (\exp_a(\tau u' + u) - \exp_a(u)) / \tau.$$

Moreover, because the cone kernel $K_{\delta t, \zeta}(a, \exp_a u)$ in (2.9) decays exponentially away from $u = 0$ for $\zeta \in [0, 1)$, it is possible to restrict the integral over \mathcal{M} in (3.1) to an integral over $B_{\delta t}$, incurring an exponentially small error as $\delta t \rightarrow 0$ (e.g., Lemma 1 in [6]). For our purposes, it suffices to consider terms up to order $\delta t^{3/2}$ in the asymptotics [15, 9], i.e.,

$$(A.3) \quad H_{\delta t, \zeta} f(a) = \frac{1}{\delta t^m} \int_{B_{\delta t}} K_{\delta t, \zeta}(a, \exp_a u) \tilde{f}(u) \det D \exp_a|_u du^1 \dots du^m + O(\delta t^{3/2}),$$

where $\tilde{f}(u) = f(\exp_a u)$. Taylor-expanding this expression about $u = 0$ using (2.10) and (2.12) with $p \geq 4$, we obtain

$$(A.4) \quad H_{\delta t, \zeta} f(a) = \frac{1}{\delta t^m} \int_{B_{\delta t}} \left(1 - \frac{h(u, u)}{\delta t^2} \right) \det D \exp_a|_u du^1 \dots du^m + O(\delta t^{3/2}),$$

where h is the Riemannian metric specified in (3.5). Note that $h = \sum_{\mu, \nu=1}^m h_{\mu\nu} U^{*\mu} U^{*\nu}$, where

$$(A.5) \quad h_{\mu\nu} = \frac{1}{\|v\|_g^2} \left(g_{\mu\nu} - \zeta \frac{v_\mu^* v_\nu^*}{\|v\|_g^2} \right), \quad v_\mu^* = \sum_{\nu=1}^m g_{\mu\nu} v^\nu,$$

and $g_{\mu\nu}$ are the components of g in (A.1).

Next, let $\widehat{\exp}_a$ and \hat{u}^μ be the exponential map and the corresponding normal coordinates associated with h , respectively. For small-enough δt , there exists $\hat{B}_{\delta t} \subset \mathbb{R}^m$ such that $\widehat{\exp}_a \hat{B}_{\delta t} = \mathcal{B}_{\delta t}$ and

$$du^1 \dots du^m = \det D \exp_a^{-1} \widehat{\exp}_a|_{\hat{u}} d\hat{u}^1 \dots d\hat{u}^m.$$

Equation (A.4) can then be expressed as

$$(A.6) \quad H_{\delta t, \zeta} f(a) = \frac{1}{\delta t^m} \int_{\hat{B}_{\delta t}} \left(1 - \frac{\|\hat{u}\|^2}{\delta t^2} \right) \hat{f}(\hat{u}) \det D \widehat{\exp}_a|_{\hat{u}} d\hat{u}^1 \dots d\hat{u}^m + O(\delta t^{3/2}),$$

with $\|\hat{u}\|^2 = \sum_{\mu=1}^m (\hat{u}^\mu)^2$ and $\hat{f}(\hat{u}) = f(\widehat{\exp}_a \hat{u})$. Following similar arguments as those used to derive (A.3), it is possible to convert (A.6) to an isotropic Gaussian integral over $\hat{B}_{\delta t}$, i.e.,

$$(A.7) \quad H_{\delta t, \zeta} f(a) = \frac{1}{\delta t^m} \int_{\hat{B}_{\delta t}} e^{-\|\hat{u}\|^2 / \delta t^2} \hat{f}(\hat{u}) \det D\widehat{\exp}_a|_{\hat{u}} d\hat{u}^1 \cdots d\hat{u}^m + O(\delta t^3).$$

It then follows from Lemma 8 of [15] (or Lemma 9 of [6]) that

$$(A.8) \quad \begin{aligned} H_{\delta t, \zeta} f(a) &= C_0 f(a) + \hat{C}_2 \Delta_{\mathbb{R}^m} \left(\hat{f}(0) \det D\widehat{\exp}_a|_{\hat{u}=0} \right) \delta t^2 + O(\delta t^3) \\ &= C_0 (f(a) + (C_2 \kappa(a) f(a) + \Delta_\zeta f(a)) \delta t^2 + O(\delta t^3)). \end{aligned}$$

Here, C_0 , C_2 , and \hat{C}_2 are constants, $\Delta_{\mathbb{R}^m}$ the canonical Laplacian in \mathbb{R}^m with $\Delta_{\mathbb{R}^m} \hat{f} = \sum_{\mu=1}^m \partial^2 \hat{f} / \partial u^{\mu^2}$, and $\kappa(a)$ a function proportional to the scalar curvature of the h metric at a .

Inserting (A.8) into (3.2), we then have that asymptotically as $\delta t \rightarrow 0$ $\mathcal{L}_{\delta t, \zeta} f(a) = C \Delta_\zeta f(a) + O(\delta t)$ for some constant C (Theorem 2.4.1 in [9]), which implies the result in (3.5).

To verify (3.6), it suffices to compute the determinants of the matrices formed by the $g_{\mu\nu}$ and $h_{\mu\nu}$ components and the vector field norm $\|v\|_g$ in any convenient coordinate system. In particular, (3.6) follows immediately by evaluating the expressions in (A.11) and (A.18) for $\|v\|_g$ and $\det h$ in terms of exponential normal coordinates for g at $u = 0$. \square

A.4. Proof of Lemma 3.2. We work in Riemannian normal coordinates u^1, \dots, u^m associated with an orthonormal basis U_1, \dots, U_m of $T_a \mathcal{M}$ such that U_1 is aligned with the dynamical vector field v at $a \in \mathcal{M}$. Thus, in this coordinate system we have

$$(A.9) \quad \begin{aligned} v^1 &= v^1(0) + \sum_{\mu=1}^m \partial_\mu v^1(0) u^\mu + O(\|u\|^2), \\ v^\nu &= \sum_{\mu=1}^m \partial_\mu v^\nu(0) u^\mu + O(\|u\|^2), \quad \nu > 1, \\ g_{\mu\nu} &= \delta_{\mu\nu} + O(\|u\|^2). \end{aligned}$$

Let $w = \sum_{\mu=1}^m w_\mu U^{*\mu}$ be a C^1 1-form field with components w_μ in the dual basis $U^{*\mu}$ to U_μ . Our objective is to evaluate $\delta_\zeta w$ from (3.3) at a using the standard expression for the codifferential in a local coordinate system (e.g., [46]),

$$(A.10) \quad \begin{aligned} \delta_\zeta w &= - \frac{1}{\sqrt{\det h}} \sum_{\mu, \nu=1}^m \partial_\mu (h^{-1, \mu\nu} \sqrt{\det h} w_\nu) \Big|_{u=0} \\ &= - \sum_{\mu, \nu=1}^m \left(h^{-1, \mu\nu} \partial_\mu w_\nu + w_\nu \partial_\mu h^{-1, \mu\nu} + \frac{h^{-1, \mu\nu} w_\nu}{2 \det h} \partial_\mu \det h \right) \Big|_{u=0}, \end{aligned}$$

and then collect the dominant terms in $(1 - \zeta)$.

First, consider the dual vector field $v^* = g(v, \cdot) = \sum_{\mu=1}^m v_\mu^* U^{*\mu}$ defined in (2.12) and its norm $\|v\|_g$ with respect to the data space metric g . It follows from (A.9) that

$$\begin{aligned} v_1^* &= v^1(0) + \sum_{\mu=1}^m \partial_\mu v^1(0) u^\mu + O(\|u\|^2), \\ v_\nu^* &= \sum_{\mu=1}^m \partial_\mu v^\nu(0) u^\mu + O(\|u\|^2), \quad \nu > 1 \\ (A.11) \quad \|v\|_g &= v^1(0) + \sum_{\mu=1}^m \partial_\mu v^1(0) u^\mu + O(\|u\|^2). \end{aligned}$$

Moreover,

$$\begin{aligned} (v_1^*)^2 / \|v\|_g^2 &= 1 + O(\|u\|^2), \\ v_1^* v_\mu^* / \|v\|_g^2 &= \frac{1}{v^1(0)} \sum_{\rho=1}^m \partial_\rho v^\mu(0) u^\rho + O(\|u\|^2), \quad \mu > 1, \\ v_\mu^* v_\nu^* / \|v\|_g^2 &= O(\|u\|^2), \quad \mu, \nu > 1. \end{aligned}$$

With these results, the metric components $h_{\mu\nu}$ in (A.5) can be expressed as

$$(A.12) \quad h_{\mu\nu} = \bar{h}_{\mu\nu} + h'_{\mu\nu} + O(\|u\|^2),$$

with

$$\begin{aligned} \bar{h}_{11} &= \frac{1 - \zeta}{(v^1(0))^2}, \quad h'_{11} = -\frac{2(1 - \zeta)}{(v^1(0))^3} \sum_{\rho=1}^m \partial_\rho v^1(0) u^\rho, \\ \bar{h}_{1\nu} &= 0, \quad h'_{1\nu} = -\frac{\zeta}{(v^1(0))^3} \sum_{\rho=1}^m \partial_\rho v^\nu(0) u^\rho, \quad \nu > 1, \\ \bar{h}_{\mu\nu} &= \frac{\delta_{\mu\nu}}{(v^1(0))^2}, \quad h'_{\mu\nu} = -\frac{2\delta_{\mu\nu}}{(v^1(0))^3} \sum_{\rho=1}^m \partial_\rho v^1(0) u^\rho, \quad \mu, \nu > 1. \end{aligned}$$

Next, consider the inverse metric $h^{-1, \mu\nu}$. Asymptotically, we have

$$(A.13) \quad h^{-1, \mu\nu} = \bar{h}^{-1, \mu\nu} - h'^{\mu\nu} + O(\|u\|^2), \quad \text{with} \quad h'^{\mu\nu} = \sum_{\rho, \sigma=1}^m \bar{h}^{-1, \mu\rho} h'_{\rho\sigma} \bar{h}^{-1, \sigma\nu}.$$

The tensor components appearing in (A.13) are

$$\begin{aligned}\bar{h}^{-1,11} &= \frac{(v^1(0))^2}{1-\zeta}, \quad h'^{11} = -\frac{2v^1(0)}{1-\zeta} \sum_{\rho=1}^m \partial_\rho v^1(0) u^\rho, \\ \bar{h}^{-1,1\nu} &= 0, \quad h'^{1\nu} = -\frac{\zeta v^1(0)}{1-\zeta} \sum_{\rho=1}^m \partial_\rho v^\nu(0) u^\rho, \quad \nu > 1, \\ \bar{h}^{-1,\mu\nu} &= (v^1(0))^2 \delta_{\mu\nu}, \quad h'^{\mu\nu} = -2v^1(0) \sum_{\rho=1}^m \partial_\rho v^1(0) u^\rho, \quad \mu, \nu > 1.\end{aligned}$$

We therefore obtain

$$(A.14) \quad - \sum_{\mu,\nu=1}^m h^{-1,\mu\nu} \partial_\mu w_\nu \Big|_{u=0} = - \sum_{\mu,\nu=1}^m \bar{h}^{-1,\mu\nu} \partial_\mu w_\nu(0) = -\frac{(v^1(0))^2}{1-\zeta} \partial_1 w_1(0) - (v^1(0))^2 \sum_{\nu=1}^m \partial_1 w_\nu(0)$$

and

$$\begin{aligned}(A.15) \quad - \sum_{\mu,\nu=1}^m \partial_\mu h^{-1,\mu\nu} w_\nu \Big|_{u=0} &= \sum_{\mu,\nu=1}^m \partial_\mu h'^{\mu\nu}(0) w_\nu(0) \\ &= -\frac{2v^1(0)}{1-\zeta} \partial_1 v^1(0) w_1(0) - \frac{\zeta v^1(0)}{1-\zeta} \sum_{\nu=2}^m (\partial_1 v^\nu(0) w_\nu(0) + \partial_\nu v^1(0) w_1(0)) \\ &\quad - 2v^1(0) \sum_{\nu=2}^m \partial_\nu v^1(0) w_\nu(0).\end{aligned}$$

What remains is to compute the contribution to the codifferential from the derivatives of $\det h$. To carry this calculation, we begin from the standard formula for the determinant,

$$(A.16) \quad \det h = \sum_{\alpha_1, \alpha_2, \dots, \alpha_m=1}^m \epsilon_{\alpha_1 \alpha_2 \dots \alpha_m} h_{1\alpha_1} h_{2\alpha_2} \cdots h_{m\alpha_m},$$

where $\epsilon_{\alpha_1 \alpha_2 \dots \alpha_m}$ is equal to 1 (-1) if $(\alpha_1, \dots, \alpha_m)$ is an even (odd) permutation of $(1, \dots, m)$, and vanishes if the α_i are not distinct. Because the $O(\|u\|)$ off-diagonal components $h_{\mu\nu}$ occur for $\mu = 1$ or $\nu = 1$ only, the $h_{i\alpha_i}$ factors

in (A.16) are $O(\|u\|)$ only for $\alpha_i \in \{1, \alpha_i\}$. Therefore,

$$\begin{aligned}
\det h &= \sum_{\alpha_1=1}^m \sum_{\alpha_2 \in \{1,2\}} \sum_{\alpha_3 \in \{1,3\}} \cdots \sum_{\alpha_m \in \{1,m\}} \epsilon_{\alpha_1 \alpha_2 \cdots \alpha_m} h_{1\alpha_1} h_{2\alpha_2} \cdots h_{m\alpha_m} + O(\|u\|^2) \\
&= \sum_{\alpha_1=1}^m \sum_{\alpha_3 \in \{1,3\}} \cdots \sum_{\alpha_m \in \{1,m\}} (\epsilon_{\alpha_1 1 \alpha_3 \cdots \alpha_m} h_{1\alpha_1} h_{21} h_{3\alpha_3} \cdots h_{m\alpha_m} \\
&\quad + \epsilon_{\alpha_1 2 \alpha_3 \cdots \alpha_m} h_{1\alpha_1} h_{22} h_{3\alpha_3} \cdots h_{m\alpha_m}) + O(\|u\|^2) \\
&= \sum_{\alpha_1=1}^m \epsilon_{\alpha_1 23 \cdots m} h_{1\alpha_1} h_{21} h_{33} \cdots h_{mm} \\
&\quad + \sum_{\alpha_1=1}^m \sum_{\alpha_3 \in \{1,3\}} \cdots \sum_{\alpha_m \in \{1,m\}} \epsilon_{\alpha_1 2 \alpha_3 \cdots \alpha_m} h_{1\alpha_1} h_{22} h_{3\alpha_3} \cdots h_{m\alpha_m} + O(\|u\|^2) \\
&= -h_{12} h_{21} h_{33} \cdots h_{mm} + \sum_{\alpha_1=1}^m \sum_{\alpha_3 \in \{1,3\}} \cdots \sum_{\alpha_m \in \{1,m\}} \epsilon_{\alpha_1 2 \alpha_3 \cdots \alpha_m} h_{1\alpha_1} h_{22} h_{3\alpha_3} \cdots h_{m\alpha_m} \\
&\quad + O(\|u\|^2).
\end{aligned}$$

Repeating this decomposition $m-2$ times, we obtain

$$\begin{aligned}
&\text{(A.17)} \\
\det h &= h_{11} h_{22} \cdots h_{mm} + (-h_{12} h_{21} h_{33} \cdots h_{mm} + h_{13} h_{22} h_{31} h_{44} \cdots h_{mm} + \cdots) + O(\|u\|^2).
\end{aligned}$$

By virtue of (A.12), the first term in the right-hand side of (A.17) is $O(\|u\|)$, whereas each of the terms in the parentheses is $O(\|u\|^2)$. Specifically,

$$\text{(A.18)} \quad \det h = \frac{1-\zeta}{(v^1(0))^{2m}} \left(1 - \frac{2}{v^1(0)} \sum_{\rho=1}^m \partial_\rho v^1(0) u^\rho \right)^m + O(\|u\|^2),$$

which implies that

$$\left. \frac{\partial_\mu \det h}{\det h} \right|_{u=0} = -\frac{2m}{v^1(0)} \partial_\mu v^1(0)$$

and

$$\begin{aligned}
&\text{(A.19)} \\
-\frac{h^{-1,\mu\nu} w_\nu}{2 \det h} \partial_\mu \det h \Big|_{u=0} &= \frac{m v^1(0)}{1-\zeta} \partial_1 v^1(0) w_1(0) + m v^1(0) \sum_{\nu=2}^m \partial_\nu v^1(0) w_\nu(0).
\end{aligned}$$

Using (A.14), (A.15), and (A.19), it is possible to write down an expression for the codifferential $\delta_\zeta w$ in the u^μ coordinate basis. For the purpose of proving Lemma 3.2, it suffices to consider the dominant, $O((1-\zeta)^{-1})$, terms, namely

$$\delta_\zeta(w) = -\frac{(v^1(0))^2}{1-\zeta} \partial_1 w_1(0) - \frac{v^1(0)}{1-\zeta} \sum_{\nu=1}^m w_\nu(0) \partial_1 v^\nu(0) - \frac{v^1(0)}{1-\zeta} \sum_{\nu=1}^m \partial_\nu v^\nu(0) w_1(0)$$

$$\begin{aligned}
&\text{(A.20)} \\
&+ \frac{m v^1(0)}{1-\zeta} \partial_1 v^1(0) w_1(0) + O((1-\zeta)^0).
\end{aligned}$$

Noting the relations

$$\begin{aligned}
w(v)|_{u=0} &= \sum_{\mu=1}^m (v^\mu w_\mu)|_{u=0} = w_1(0)v^1(0), \\
v(w(v))|_{u=0} &= \sum_{\mu,\nu=1}^m v^\mu \partial_\mu (w_\nu v^\nu)|_{u=0} = \sum_{\mu,\nu=1}^m (v^\mu v^\nu \partial_\mu w_\nu + v^\mu w_\nu \partial_\mu v^\nu)|_{u=0} \\
&= (v^1(0))^2 \partial_1 w_1(0) + v^1(0) \sum_{\nu=1}^m v^1(0) w_\nu(0) \partial_1 v^\nu(0), \\
\operatorname{div}_\nu v|_{u=0} &= \frac{1}{\sqrt{\det h}} \sum_{\nu=1}^m \partial_\nu (\sqrt{\det h} v^\nu) \Big|_{u=0} = \sum_{\nu=1}^m \partial_\nu v^\nu(0) - m \partial_1 v^1(0), \\
\operatorname{div}_\nu [w(v)v] &= w(v) \operatorname{div}_\nu v + v(w(v)),
\end{aligned}$$

the asymptotic expansion (A.20) can be put in the covariant (basis-independent) form (3.7), thus proving the Lemma. \square

Acknowledgments. The author would like to thank T. Berry, M. Bushuk, J. Harlim, and A. Majda for stimulating discussions. This work was supported by ONR DRI grant N00014-14-1-0150 and ONR MURI grant 25-74200-F7112.

REFERENCES

- [1] M. A. ALEXANDER, C. DESER, AND M. S. TIMLIN, *The reemergence of SST anomalies in the North Pacific Ocean*, J. Climate, 12 (1999), pp. 2419–2433.
- [2] L. ARNOLD, *Random Dynamical Systems*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 1998.
- [3] N. AUBRY, R. GUYONNET, AND R. LIMA, *Spatiotemporal analysis of complex signals: Theory and applications*, J. Stat. Phys., 64 (1991), pp. 683–739.
- [4] N. AUBRY, W.-Y. LIAN, AND E. S. TITI, *Preserving symmetries in the proper orthogonal decomposition*, SIAM J. Sci. Comput., 14 (1993), pp. 483–505.
- [5] M. BELKIN AND P. NIYOGI, *Laplacian eigenmaps for dimensionality reduction and data representation*, Neural Comput., 15 (2003), pp. 1373–1396.
- [6] ———, *Towards a theoretical foundation for Laplacian-based manifold methods*, J. Comput. Syst. Sci., 74 (2008), pp. 1289–1308.
- [7] M. BELKIN, P. NIYOGI, AND V. SINDHWANI, *Manifold regularization: A geometric framework for learning from labeled and unlabeled examples*, J. Mach. Learn. Res., 7 (2006), pp. 2399–2434.
- [8] P. H. BÉRARD, *Embedding Riemannian manifolds by their heat kernel*, Geom. Funct. Anal., 4 (1984), pp. 373–398.
- [9] T. BERRY, *Model Free Techniques for Reduction of High-Dimensional Dynamics*, PhD thesis, George Mason University, 2013.
- [10] T. BERRY, R. CRESSMAN, Z. GREGURIC FERENČEK, AND T. SAUER, *Time-scale separation from diffusion-mapped delay coordinates*, SIAM J. Appl. Dyn. Sys., 12 (2013), pp. 618–649.
- [11] D. S. BROOMHEAD AND G. P. KING, *Extracting qualitative dynamics from experimental data*, Phys. D, 20 (1986), pp. 217–236.
- [12] M. BUSHUK, D. GIANNAKIS, AND A. J. MAJDA, *Reemergence mechanisms for North Pacific sea ice revealed through nonlinear Laplacian spectral analysis*, J. Climate, (2013). submitted.

- [13] A. CAPONNETTO, C. A. MICCHELLI, M. PONTIL, AND Y. YING, *Universal multi-task kernels*, J. Mach. Learn. Res., 9 (2008), pp. 1615–1646.
- [14] R. R. COIFMAN ET AL., *Geometric diffusions as a tool for harmonic analysis and structure definition on data*, Proc. Natl. Acad. Sci., 102 (2005), pp. 7426–7431.
- [15] R. R. COIFMAN AND S. LAFON, *Diffusion maps*, Appl. Comput. Harmon. Anal., 21 (2006), pp. 5–30.
- [16] ———, *Geometric harmonics: A novel tool for multiscale out-of-sample extension of empirical functions*, Appl. Comput. Harmon. Anal., 21 (2006), pp. 31–52.
- [17] W. D. COLLINS ET AL., *The Community Climate System Model version 3 (CCSM3)*, J. Climate, 19 (2006), pp. 2122–2143.
- [18] D. T. CROMMELIN AND A. J. MAJDA, *Strategies for model reduction: Comparing different optimal bases*, J. Atmos. Sci., 61 (2004), pp. 2206–2217.
- [19] M. DESBRUN, E. KANSO, AND Y. TONG, *Discrete differential forms for computational modeling*, in ACM SIGGRAPH 2005 Courses, SIGGRAPH '05, Los Angeles, 2005.
- [20] P. DEUFLHARD, M. DELLNITZ, O. JUNGE, AND C. SCHÜTTE, *Computation of essential molecular dynamics by subdivision techniques I: basic concept*, Lect. Notes Comp. Sci. Eng., 4 (1999), p. 98.
- [21] E. DI LORENZO ET AL., *North Pacific Gyre Oscillation links ocean climate and ecosystem change*, Geophys. Res. Lett., 35 (2008), p. L08607.
- [22] V. P. DYMNIKOV AND A. N. FILATOV, *Mathematics of Climate Modeling*, Birkhäuser, Boston, 1997.
- [23] K. D. ELWORTHY, Y. LE JAN, AND X.-M. XI, *The Geometry of Filtering*, Frontiers in Mathematics, Birkhäuser, 2010.
- [24] M. GHIL ET AL., *Advanced spectral methods for climatic time series*, Rev. Geophys., 40 (2002).
- [25] D. GIANNAKIS AND A. J. MAJDA, *Time series reconstruction via machine learning: Revealing decadal variability and intermittency in the North Pacific sector of a coupled climate model*, in Conference on Intelligent Data Understanding 2011, Mountain View, California, 2011.
- [26] ———, *Comparing low-frequency and intermittent variability in comprehensive climate models through nonlinear Laplacian spectral analysis*, Geophys. Res. Lett., 39 (2012), p. L10710.
- [27] ———, *Limits of predictability in the North Pacific sector of a comprehensive climate model*, Geophys. Res. Lett., 39 (2012), p. L24602.
- [28] ———, *Nonlinear Laplacian spectral analysis for time series with intermittency and low-frequency variability*, Proc. Natl. Acad. Sci., 109 (2012), pp. 2222–2227.
- [29] ———, *Nonlinear Laplacian spectral analysis: Capturing intermittent and low-frequency spatiotemporal patterns in high-dimensional data*, Stat. Anal. Data Min., 6 (2013), pp. 180–194.
- [30] L. G. GRADY AND J. R. POLIMENI, *Discrete Calculus*, Springer-Verlag, London, 2010.
- [31] M. HEIN, J.-Y. AUDIBERT, AND U. VON LUXBURG, *From graphs to manifolds—weak and strong pointwise consistency of graph Laplacians*, in Learning Theory, vol. 3559 of Lecture Notes in Computer Science, Springer, Berlin, 2005, pp. 470–485.
- [32] P. HOLMES, J. L. LUMLEY, AND G. BERKOOZ, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, 1996.
- [33] P. W. JONES, M. MAGGIONI, AND R. SCHUL, *Manifold parametrizations by eigenfunctions of the Laplacian and heat kernels*, Proc. Natl. Acad. Sci., 105 (2008), pp. 1803–1808.
- [34] A. KATOK AND H. B., *Introduction to the Modern Theory of Dynamical Systems*, Cambridge University Press, Cambridge, 1997.
- [35] M. A. KATSOUKAKIS, A. J. MAJDA, AND A. SOPSAKIS, *Intermittency, metastability and coarse-graining for coupled deterministic–stochastic lattice systems*, Nonlinearity, 19 (2006), pp. 1021–1047.
- [36] M. A. KATSOUKAKIS AND D. G. VLACHOS, *Coarse-grained stochastic processes and*

- kinetic Monte Carlo simulators for the diffusion of interacting particles*, J. Chem. Phys., 119 (2003), pp. 9412–9427.
- [37] J. A. LEE AND M. VERLEYSEN, *Nonlinear Dimensionality Reduction*, Information Science and Statistics, Springer, New York, 2007.
 - [38] R. J. LEVEQUE, *Finite Difference Methods for Ordinary and Partial Differential Equations*, Classics in Applied Mathematics, Society for Industrial and Applied Mathematics, Philadelphia, 2007.
 - [39] R. A. MADDEN AND P. R. JULIAN, *Description of global-scale circulation cells in the tropics with a 40–50 day period*, J. Atmos. Sci., 29 (1972), pp. 1109–1123.
 - [40] A. J. MAJDA AND X. WANG, *Nonlinear Dynamics and Statistical Theories for Basic Geophysical Flows*, Cambridge University Press, Cambridge, 2006.
 - [41] N. J. MANTUA AND S. R. HARE, *The Pacific Decadal Oscillation*, J. Oceanogr., 58 (2002), pp. 35–44.
 - [42] C. A. MICCHELLI AND M. PONTIL, *On learning vector-valued functions*, Neural Comput., 17 (2005), pp. 177–204.
 - [43] B. NADLER, S. LAFON, R. R. COIFMAN, AND I. KEVRIKEDES, *Diffusion maps, spectral clustering, and reaction coordinates of dynamical systems*, Appl. Comput. Harmon. Anal., 21 (2006), pp. 113–127.
 - [44] N. H. PACKARD ET AL., *Geometry from a time series*, Phys. Rev. Lett., 45 (1980), pp. 712–716.
 - [45] J. PORTEGIES, *Embeddings of Riemannian manifolds with heat kernels and eigenfunctions*. Working paper, Courant Institute of Mathematical Sciences.
 - [46] S. ROSENBERG, *The Laplacian on a Riemannian Manifold*, vol. 31 of London Mathematical Society Student Texts, Cambridge University Press, Cambridge, 1997.
 - [47] T. SAUER, J. A. YORKE, AND M. CASDAGLI, *Embedology*, J. Stat. Phys., 65 (1991), pp. 579–616.
 - [48] B. SCHÖLKOPF, A. SMOLA, AND K. MÜLLER, *Nonlinear component analysis as a kernel eigenvalue problem*, Neural Comput., 10 (1998), pp. 1299–1319.
 - [49] A. SINGER, *From graph to manifold Laplacian: The convergence rate*, J. Appl. Comput. Harmon. Anal., 21 (2006), pp. 128–134.
 - [50] K. R. SREENIVASAN AND R. A. ANTONIA, *The phenomenology of small-scale turbulence*, Annu. Rev. Fluid Mech., 29 (1997), pp. 435–472.
 - [51] F. TAKENS, *Detecting strange attractors in turbulence*, in Dynamical Systems and Turbulence, Warwick 1980, vol. 898 of Lecture Notes in Mathematics, Springer, Berlin, 1981, pp. 366–381.
 - [52] R. TALMON AND R. R. COIFMAN, *Empirical intrinsic geometry for nonlinear modeling and time series filtering*, Proc. Natl. Acad. Sci., 110 (2013), pp. 12535–12540.
 - [53] K. E. TRENBERTH, *The definition of El Niño*, Bull. Amer. Meteor. Soc., 78 (1997), pp. 2771–2777.
 - [54] R. VAUTARD AND M. GHIL, *Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series*, Phys. D, 35 (1989), pp. 395–424.
 - [55] D. ZHOU AND C. BURGESS, *High-order regularization on graphs*, in 6th International Workshop on Mining and Learning with Graphs, Helsinki, 2008.